

Department of the Navy Information Technology Infrastructure Architecture

Attachment B

- *Executive Summary*
- *Network Infrastructure and Services Architecture (ITIA Volume I)*
- *Enterprise Architecture Framework (ITIA Volume II)*

Department of the Navy Chief Information Officer (DON CIO)
Information Technology Infrastructure Integrated Product Team (ITI IPT)

Version 1.0 proposed
16 March 1999

This page intentionally left blank.

Department of the Navy Information Technology Infrastructure Architecture

Executive Summary

Department of the Navy Chief Information Officer (DON CIO)
Information Technology Infrastructure Integrated Product Team (ITI IPT)

Version 1.0 proposed
16 March 1999

This page intentionally left blank.

Table of Contents

| | |
|--|---|
| 1. Background and Purpose | 1 |
| 2. ITI Architecture Strategic Drivers | 1 |
| 3. ITI Architecture Document Overview | 2 |
| Volume I - Network Infrastructure and Service Architecture | 3 |
| Volume II - Enterprise Architecture Framework..... | 5 |
| 4. Governance | 7 |
| 5. Implementation Recommendations | 8 |

This page intentionally left blank.

1. Background and Purpose

In May 1998, representatives from major echelon 2 Department of the Navy (DON) organizations met in the forum of the DON CIO Board of Representatives to discuss the need for an enterprise information technology infrastructure. As a result of that meeting, the DON Information Technology Infrastructure (ITI) Integrated Product Team (IPT) was chartered by the Board of Representatives. The ITI IPT was tasked with developing the DON information technology infrastructure architecture framework, guidance, and templates for all DON ITI network connectivity and basic network services, including end-to-end network management and security. In December 1998, the ITI IPT released the draft of the ITI Architecture to the Board of Representatives for review. The final version incorporates the resulting comments and recommendations.

This Executive Summary provides an overview of the ITI IPT report. For more information, see the two published volumes of the report contained in this document and accessible through the DON CIO website at <http://www.doncio.navy.mil/>.

The DON ITI Architecture is targeted at providing a DON network of networks that meets the ITI security, interoperability, and performance requirements of the Navy and Marine Corps, including both tactical and non-tactical missions. The ITI Architecture fully supports and provides implementation guidance for approved DON strategies, including IT-21 and Network Centric Warfare, and is aligned with the Department of Defense (DoD) C4ISR Framework 2.0 and Joint Technical Architecture (2.0).

The ITI Architecture, along with its companion document, the Information Technology Standards Guidance (ITSG) Version 98-1.1 (released in June 1998), enables network planners and service providers to design, develop, and implement integrated network solutions that are seamless and cost effective. The ITI Architecture is designed to support the DON's ITI management philosophy of centralized IT policy and decentralized implementation. All DON personnel involved in network planning, design, and procurement should use the ITI Architecture and the ITSG documents to formulate integrated plans and to undertake IT modernization initiatives.

2. ITI Architecture Strategic Drivers

An ITI Architecture provides the foundation for robust DON enterprise-wide ITI. Four major strategic forces are driving the need for this robust infrastructure.

- Dependency on ITI to achieve Naval information superiority and the potential of ITI to support a revolution in military affairs (RMA) and revolution in business affairs (RBA).
- Alignment of ITI investment initiatives to produce a more focused, efficient, and holistic approach to building and operating the DON enterprise infrastructure.
- The need to correct the current imbalance in resources applied to combat programs versus combat support programs.
- The requirement to implement the Clinger-Cohen Act of 1996 and Office of Management and Budget Memorandum 97-16.

3. ITI Architecture Document Overview

The ITI IPT produced a DON ITI Architecture document that is organized into multiple volumes.

- Volume I - Network Infrastructure and Services Architecture, describes the network connectivity and network services architecture by components. It also provides architecture guidance in the form of a Wide Area Network Connectivity Plan, the Metropolitan Area Network Design Template, the Campus Area Network Design Template, and an Information Technology Services Center (ITSC) template.
- Volume II - Enterprise Architecture Framework, introduces a comprehensive structure for integrating all major planning activities involved in the transformation of the DON. It positions the critical role of IT infrastructure in supporting the evolving operational requirements of the DON within the context of identified strategic mission requirements. A framework for the ITI Architecture is also provided.
- There are a number of additional components of the ITI Infrastructure Architecture that remain to be developed. These will be added to this document structure as they are completed and approved. In addition, the ITI IPT presented a number of issues and recommendations relating to ITI governance in the DON. These presentations include strong recommendations for changes in policy, organization, and business processes with respect to ITI architecture, requirements planning, design, procurement, management, and operations. As decisions are made regarding changes in ITI governance, it is our intention to publish them as extensions or additional volumes to this architecture document.

Volume I - Network Infrastructure and Service Architecture

Volume I - Network Infrastructure and Services Architecture provides architecture guidance to address the enterprise network requirements of the DON. It identifies the special DON geographical, organizational, and operating environments that must be accommodated to satisfy the diversity of ITI users. The conceptual network architecture for the DON is represented in Figure 1.

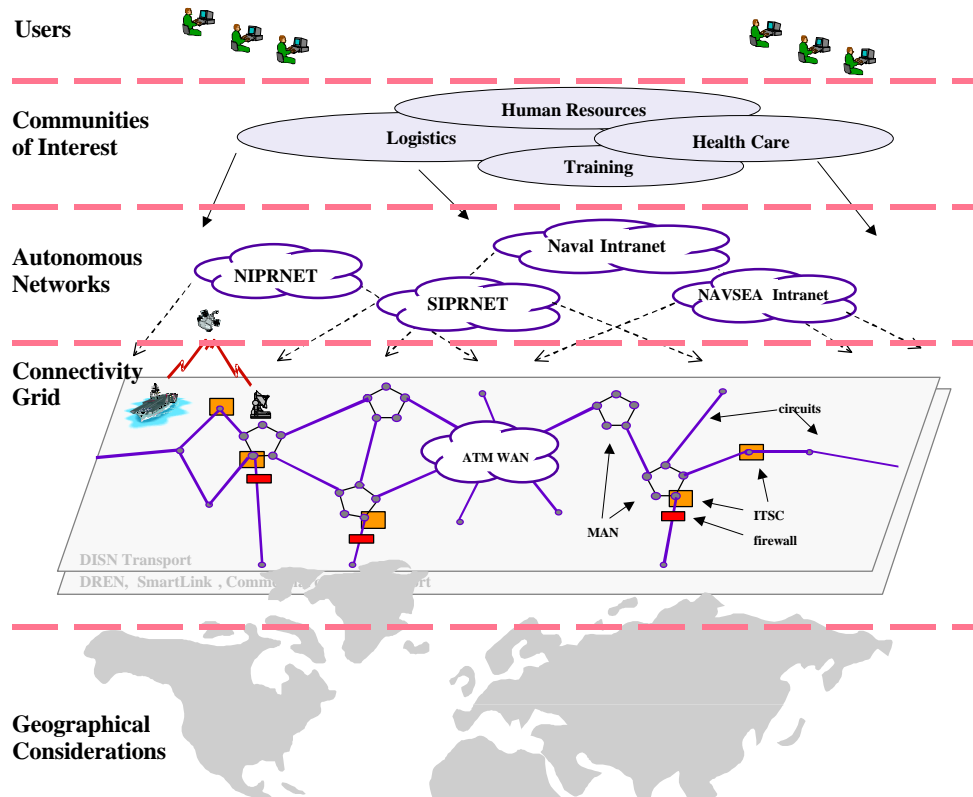


Figure 1. Conceptual Network Architecture

The network infrastructure should align geographically with Naval user population densities providing global end-to-end connectivity. The underlying connectivity grid provides the transport capability for all network services. Autonomous networks include the specific service networks (e.g., SIPRNET) that provide capabilities targeted at specific requirements. Communities of interest represent the numerous functional areas (e.g., medical health system) and their associated information management and systems applications that must be network accessible by the users working in these functional areas. These users are at geographically and organizationally dispersed sites, and many of the users are members of multiple communities of interest. Collectively, these enterprise networks must meet the diverse requirements of all Navy and Marine Corps users – for seamlessness, mobility, and jointness.

Translating the conceptual architecture to a network solution is facilitated by the technical model depicted in Figure 2. The model was adopted as a basis for considering, organizing, developing, and presenting the components of this ITI architecture. The network-centric solution begins at the bottom layer with the transmission service. Above this layer are the additional layers (e.g., SONET rings, ATM networks, and IP Networks) of networking capability that successfully interact according to the rules of this architecture to deliver the required network services.

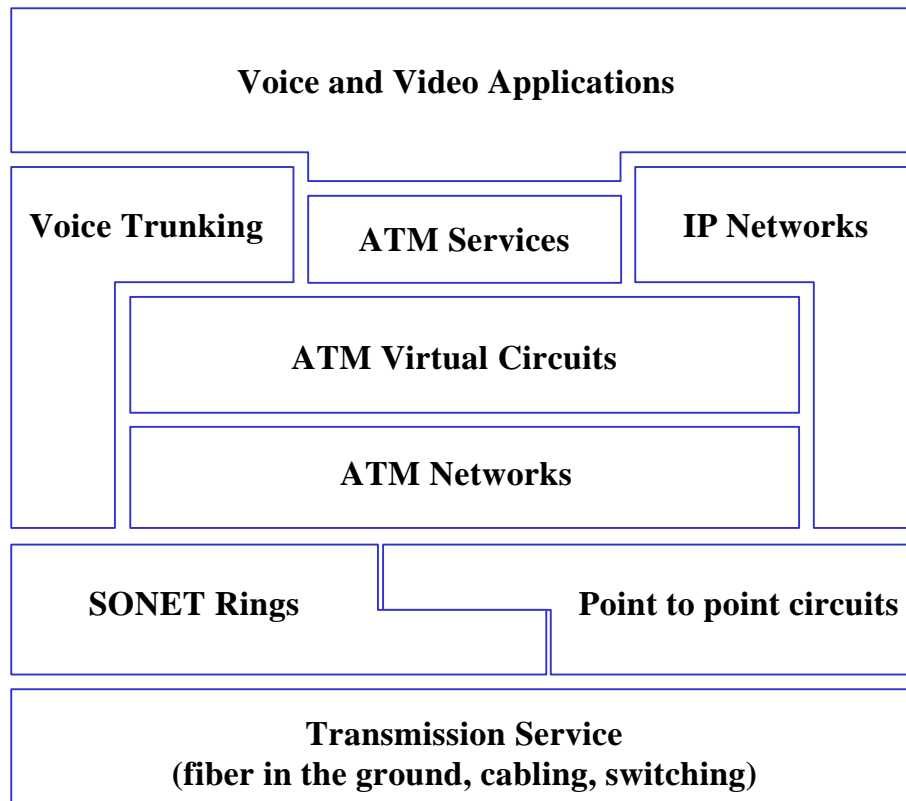


Figure 2. Layered Approach to Network Solution

Highlighting Volume I, it contains detailed guidance for deployment of an ATM infrastructure across the DON's Wide Area Networks (WANs) and Metropolitan Area Networks (MANs), and to the campuses. It fully supports the DON's ubiquitous IP networks and, in fact, provides improved performance through IP over ATM. The architecture also provides details of essential network services required for interconnectivity and operation of the DON enterprise network. It includes the following:

- A robust and responsive ATM solution in the WAN and MAN environments that recognizes multiple ATM service implementations, supports a wide range of DON service level requirements, and provides a solid capability to support virtual private networks.
- A full description of the enterprise solution for deploying internal and external routing protocols to provide efficient interconnectivity and improved network performance.
- Specific security solutions for the ATM and IP network environments by taking the established DoD principles of Defense in Depth to the next level of detail.
- Detailed description of the 14 essential enterprise network services that allow the design of DON enterprise solutions to permit efficient network component interconnectivity and interoperability. Examples of enterprise services are Directory Services, Domain Name System (DNS), and Network Time Service.
- Support requirements as defined by IT-21, including point-to-point ATM connectivity to the desktop for those users that require it.

- A Wide-Area Connectivity Plan for a single enterprise ATM backbone service to meet global enterprise-wide needs and to provide the end-to-end connectivity required for DON functional missions.
- The connectivity guidance appropriate for a typical Naval concentration area through a detailed MAN template, and similarly, the extensive connectivity guidance for bases presented in a very detailed CAN template. The campus template can be extended to address shipboard LAN requirements.
- A Network Security Strategy that includes implementation guidelines for a system of complementary and redundant Defense In Depth network security mechanisms.

Volume II - Enterprise Architecture Framework

Volume II – Enterprise Architecture Framework introduces a comprehensive structure for supporting the integration of all major planning activities involved in the transformation of the DON. The Enterprise Architecture Framework (EAF) was developed to position the critical role of IT Infrastructure in supporting the evolving operational requirements of DON within the context of identified strategic mission requirements.

The EAF is, therefore, more than an integrated IT architecture framework. It pulls together, in one unified structure, all of the components of the enterprise that must be considered in establishing readiness to perform the DON's missions. However, the EAF strongly and properly positions the role of information (knowledge), information management systems, and the underlying IT infrastructure by linking these elements to mission essential tasks and the unique operational environments of the Navy and Marine Corps. Consequently, it is ideally suited for supporting the complex planning of RMA/RBA and giving substance to the key strategic notions of network-centric warfare and information superiority.

The high level view of the EAF is show in Figure 3.

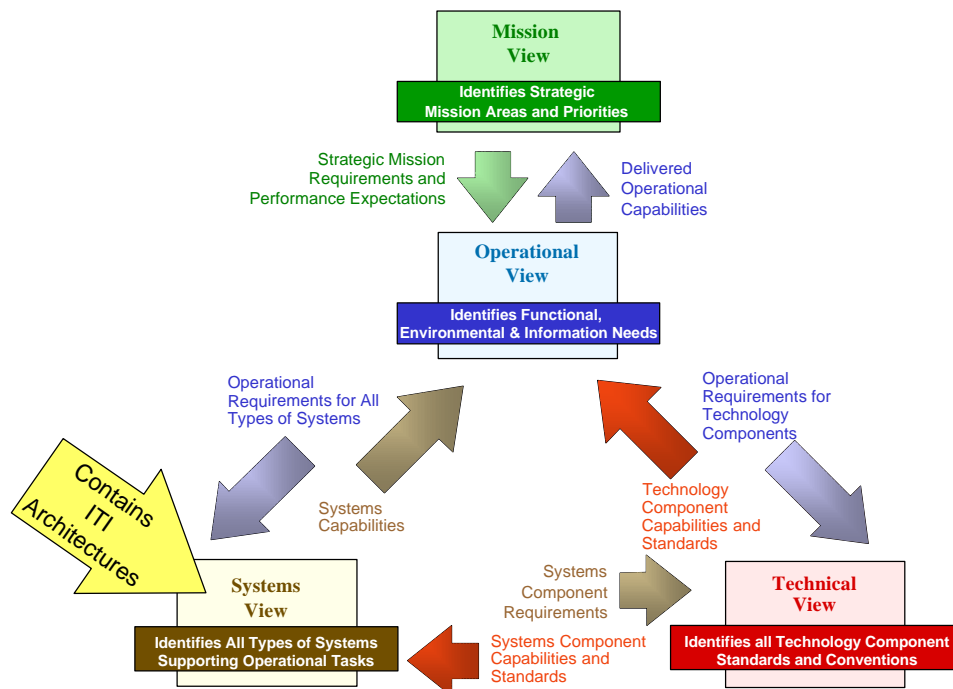


Figure 3. Enterprise Architecture Framework

The DON EAF is upwardly compatible with the existing C4ISR V2 Architecture Framework from DoD. It builds upon this established base in the following areas:

- Extends the functional task model beyond C4ISR to include all support functions as well as enterprise planning and policy setting activities so that it is a complete enterprise functional model.
- Adds a “Mission View” to the framework to provide structure and linkage to the analysis of threats, determination of strategic mission requirements, and positioning of key roles of other services, agencies, and contractors.
- Expands the “Systems View” to include all types of systems (sensor, weapon, platforms, information systems, etc.) that collectively support the warfighter and all other personnel performing mission essential tasks. The IT infrastructure is positioned within this view as a capability upon which these other systems depend.
- Provides the construct for organizing the planning of the ITI architecture by defining the sub-architecture components of Workgroup Computing, Network, and Enterprise Server combined with the required end-to-end services of Security Management, Infrastructure Management, and Information Distribution.

The ITI architecture framework is shown below in Figure 4.

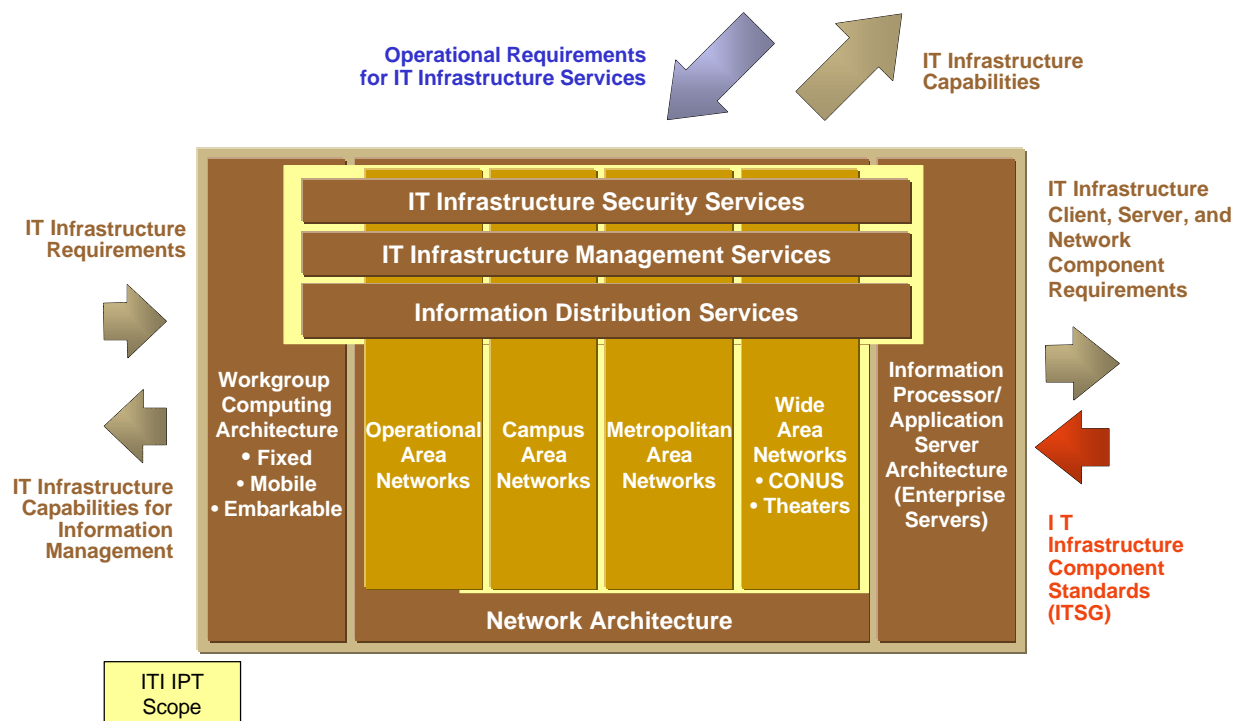


Figure 4. ITI Architecture Framework

The development of the enterprise network architecture templates and plans as presented in Volume I represents the initial application of the EAF. It will continue to be applied for further expansion of the ITI and for work on enterprise systems and enterprise processes. The broader the use of the EAF in providing structure and context for integrating enterprise planning initiatives, the greater its value to the DON and, through extension, to DoD and other agencies.

4. Governance Issues

The publication of the network architecture plans and templates addresses the technical challenge of integrating the disparate networking solutions that exist today into an effective consolidated enterprise infrastructure. The shift to an infrastructure orientation implies that many existing enterprise policies, organizational responsibilities, and business processes will be impacted by this approach.

The most critical success factor for building the DON ITI infrastructure is implementing an effective approach for ITI governance. The ITI IPT identified a number of governance issues. Five are presented here because of their importance.

- How should we organize within the DON to effectively perform the management functions of ITI? There are currently multiple organizations performing duplicate, uncoordinated functions, including planning, designing, building, provisioning, and managing networks. How should we reorganize to more effectively allocate resources, assign responsibilities, instill discipline, and maintain operational support for an enterprise network infrastructure?
- How should we pay for ITI services? Some common set of ITI services such as telephone, email, directory, time, domain naming, web access, and others must be standardized and offered as basic network services across the DON. A single DON enterprise ATM backbone and regional MANs are also network services that should be centrally funded. Any solution must incorporate mechanisms that curb network user appetites and discipline its use. How should we fund these services?
- How do we collaborate to ensure that all user requirements are reflected in enterprise solutions? How do we ensure that the regional and enterprise providers of IT infrastructure services are responsive, efficient, and accountable to their customers? How can we ensure that the relative priority placed on competing services accurately reflects the relative priority of the Navy and Marine Corps missions being supported by the ITI? How do we ensure that service providers incorporate best practices and remain cost competitive?
- The reorganization of the management functions of ITI cannot succeed without a reengineering of supporting ITI enterprise processes such as determining requirements, negotiating service agreements, reporting performance, managing operations, providing user support, and billing customers. How do we undertake an effective reengineering of our ITI processes?
- How do we address the acquisition and development of the necessary skills and competencies to effectively plan, implement, and operate this ITI? What is the best use of contractors and external service providers in meeting these capabilities?

5. Implementation Recommendations

The following initial actions are recommended to address implementation of the ITI network solution.

- Adopt the use of the EAF and the ITI architecture templates for all related enterprise planning activities.
- Determine organizational accountabilities for planning, procuring, and operating the ITI.
- Determine a Navy and Marine Corps solution for a Wide Area Connectivity Plan that results in a single DON enterprise backbone.
- Define the regions (MANs) and an operational management structure that is acceptable to the fleets, Marine Corps, and the Systems Commands and specifies responsibilities and accountability (including Information Technology Support Centers).
- Direct the full architecture development of ITI network services (directory, DNS, time, etc.) and implement them across the DON.
- Determine a funding structure for common user services and enterprise networks, and develop a billing structure for other uses.
- Develop the required ITI management processes.
- Fully provision the DON enterprise backbone and begin to connect organizations and regions.
- Conduct a pilot MAN and ITSC implementation using the ITI architecture guidance.

Implementing the IT infrastructure is one of the most strategic initiatives for the DON. It will require effective collaboration and leadership to develop and maintain this increased capability. The publication of the network architecture plans and templates is an important first step. Implementation efforts to apply this guidance are already underway. It will take the collective energy of all involved in IT to quickly and efficiently migrate to this infrastructure approach.

Department of the Navy Information Technology Infrastructure Architecture

Volume I *Network Infrastructure and* *Services Architecture*

Department of the Navy Chief Information Officer (DON CIO)
Information Technology Infrastructure Integrated Product Team (ITI IPT)

Version 1.0 proposed
16 March 1999

This page intentionally left blank.

Volume I, Chapter 1 – Table of Contents

| | |
|---|-------------------------------------|
| Acknowledgements | Error! Bookmark not defined. |
| 1. Introduction | 1-1 |
| 1.1 Document Purpose..... | 1-1 |
| 1.2 Document Structure..... | 1-1 |
| 1.3 Strategic Drivers..... | 1-2 |
| 1.4 Positioning the ITI in the Enterprise Architecture Framework..... | 1-3 |
| 1.4.1 Introduction to Enterprise Architecture Framework..... | 1-3 |
| 1.4.2 From Framework to ITI Architecture | 1-7 |

This page intentionally left blank.

1. Introduction

1.1 Document Purpose

This document defines the Department of the Navy (DON) network architecture. The network architecture is part of the DON Information Technology Infrastructure (ITI) Architecture. It provides the basis to produce a secure, interoperable, and robust DON enterprise networking capability. This document, Volume I - Network Infrastructure and Services Architecture, describes the network connectivity and network services, the Wide Area Network Connectivity Plan, the Metropolitan Area Network Design Template, the Campus Area Network Design Template, and an Information Technology Services Center (ITSC) Template.

Volume II- Enterprise Architecture Framework, is a separate publication which identifies the overarching comprehensive set of models that positions the DON ITI Architecture to support the operational requirements of DON. It further identifies the additional components of the ITI, specifically, the Workgroup Computing Architecture and the Server Architecture. The network infrastructure and services architecture described in this document are the first instantiations of the framework. Other architecture components will be developed in subsequent efforts.

This document provides the guidance for planning, developing, implementing, and operating all activities associated with DON IT network infrastructure. It is to be used by DON acquisition programs, organizations, working groups, and Integrated Product Teams (IPTs) to facilitate convergence on a single, comprehensive ITI architecture. This guidance and associated design templates are not intended to be detailed design and implementation plans, but to serve as frames of reference for design and implementation efforts.

1.2 Document Structure

Volume I consists of four chapters that describe the components and relationships of the building blocks of the enterprise architecture. It also provides the strategy and guidance for implementing wide area, metropolitan, and campus networks and a series of architecture templates that provide the planners and implementers with a tool for actual implementation. Specifically, they are as follows:

Chapter 1 places the ITI Volume I into the context of the overarching Enterprise Architecture Framework (EAF). It introduces the components that represent the building blocks of the enterprise architecture.

Chapter 2 describes the aggregation of DON network infrastructure requirements, including the large diverse group of tactical and functional users, mission systems, multiple operational areas, and all types of media.

Chapter 3 contains four major subsections:

- It introduces the basic planning construct for a global enterprise vision of the DON network connectivity architecture – the underlying connectivity grid, the autonomous networks supporting specific functions or technologies or security needs, and the communities of interest (CoI) used to logically group user processes and requirements.

- It provides a high-level overview of the DON ITI connectivity architecture that addresses both Asynchronous Transfer Mode (ATM) and Internet Protocol (IP), their individual strategies, and their interrelationships.
- It addresses the detailed ATM connectivity architecture. ATM technology forms the principal technology for the enterprise architecture. This enables the consolidation of voice, video, and data onto a single network and offers the advantages of scalability, speed, quality of service, economy, and interoperability. Included here are the pertinent protocols and references that are essential for successful ATM implementation.
- It addresses the detailed Internet Protocol (IP) connectivity architecture. IP-based connectivity components comprise a significant portion of the ITI. Support of IP in the ATM environment is an area of major focus. A number of protocols and references which are necessary to support IP networking are outlined here.

Chapter 4 provides a description of the basic network services that users require in all functional areas to be accessible from the network. All network services must have a common planning framework and consistent implementation strategy. Some services must be implemented under an enterprise plan.

There are a series of five networking-related appendices that are designed for use by planners for specific implementations. There is also an appendix that acknowledges the contributors to this document and provides point of contacts for particular subject areas.

Appendix A outlines the strategy, planning factors, and steps required to build a DON WAN.

Appendix B presents the design template required to plan a MAN, including supporting protocols, security, and management.

Appendix C presents the design template required to plan a campus network, including supporting protocols, security, and management. The term campus area network includes shipboard networks. Differences from shipboard and terrestrial campus are shown by exception.

Appendix D provides detailed guidance on the functions and processes that will be performed by the ITSC. A series of ITSCs will support shipboard and terrestrial requirements for voice, video, and data.

Appendix E provides an overview of the network security Defense in Depth model and describes the mechanisms that provide protection at the four layers.

Appendix F lists the contributors to this document.

1.3 Strategic Drivers

An ITI Architecture provides the foundation for robust DON enterprise-wide ITI. Four major strategic forces are driving the need for this robust infrastructure.

- Dependency on ITI to achieve Naval information superiority and the potential of ITI to support a revolution in military affairs (RMA) and revolution in business affairs (RBA).

- Alignment of ITI investment initiatives to produce a more focused, efficient, and holistic approach to building and operating the DON enterprise infrastructure.
- The need to correct the current imbalance in resources applied to combat programs versus combat support programs.
- The requirement to implement the Clinger-Cohen Act of 1996 and Office of Management and Budget Memorandum 97-16.

1.4 Positioning the ITI in the Enterprise Architecture Framework

The Enterprise Architecture Framework (EAF), introduced in Volume II, provides the context for describing a DON enterprise architecture through a unified and common set of reference models. These models provide a basis for describing a design target and coordinating changes across the enterprise infrastructure, the functional areas, and the joint/allied communities.

1.4.1 Introduction to Enterprise Architecture Framework

The DON is a complex, multi-business industry that combines the major elements of warfighting with a wide array of combat support activities, including research and development, engineering, transportation, acquisition, distribution, meteorology, education, healthcare, hospitality, and municipal services. The transformation of the IT services that support these functions requires a shift in DON IT strategic planning to a more interoperable, flexible, and secure IT infrastructure, more leveraging of emerging resources, improved management and communication systems, and new ways to support the warfighter with much fewer resources.

The management, organizational, and technological challenges that must be negotiated to produce a transformational shift are dynamic, complex, and intertwined. The EAF provides an overall structure, elements, and concepts to plan, prioritize opportunities, integrate plans, coordinate implementation strategies, and communicate changes.

Key portions of the EAF are the templates that support various ITI planning and design across DON communities. A template is a pre-structured approach for defining and collecting the information used by a planning or design activity. The templates enable capture of recognized commonalities in requirements and solutions across the enterprise. They provide widely-applicable architecture designs that enable reduction of planning efforts, opportunities to leverage or share solutions, and shorter implementation periods. By adopting a common framework to integrate all of the models and components related to ITI planning, various DON communities can better communicate their requirements and ideas. The end result is an EAF populated with DON-specific components that supports many ongoing planning activities for enterprise solutions.

Figure 1-1 summarizes the transformation from vision to operational capability as supported by the architecture models and templates.

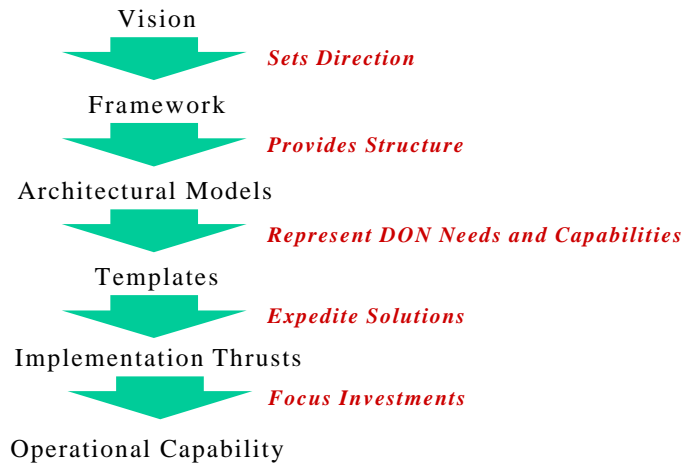


Figure 1-1. Enterprise Architecture Framework Process

The DoD C4ISR Framework, Version 2.0, which provides the basis for the DON EAF, consists of three different architecture views - Operational, Systems, and Technical. The DON EAF introduces a fourth – Mission view – which defines the strategic mission requirement. This is depicted in Figure 1-2 and is fully described in Volume II.

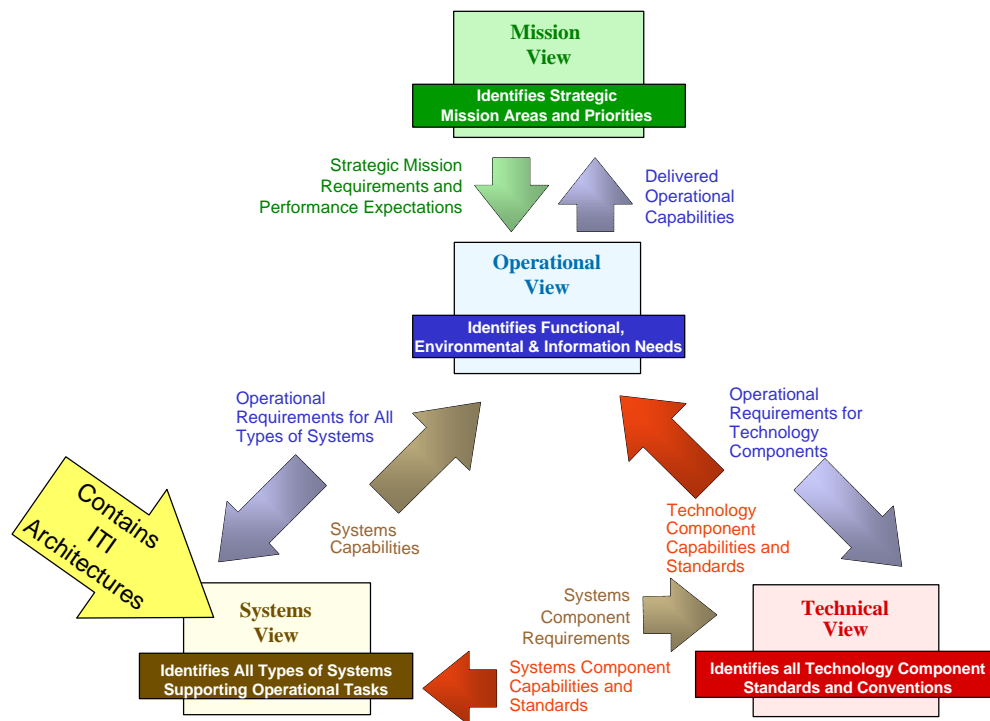


Figure 1-2. DON Architecture Framework

Each of these four views has an inter-related set of requirements and capabilities flows that interact to produce the required set of enterprise capabilities. All four views are important in determining and describing information requirements; matching those to an appropriate mix of technology solutions; and performing the required planning, design, and integration of the IT Infrastructure (ITI). The Systems

View is most relevant to the ITI architecture. By “drilling down” in the Systems View, a more detailed view of the underlying subsystems is revealed. Figure 1-3 shows this Sub-Systems View.

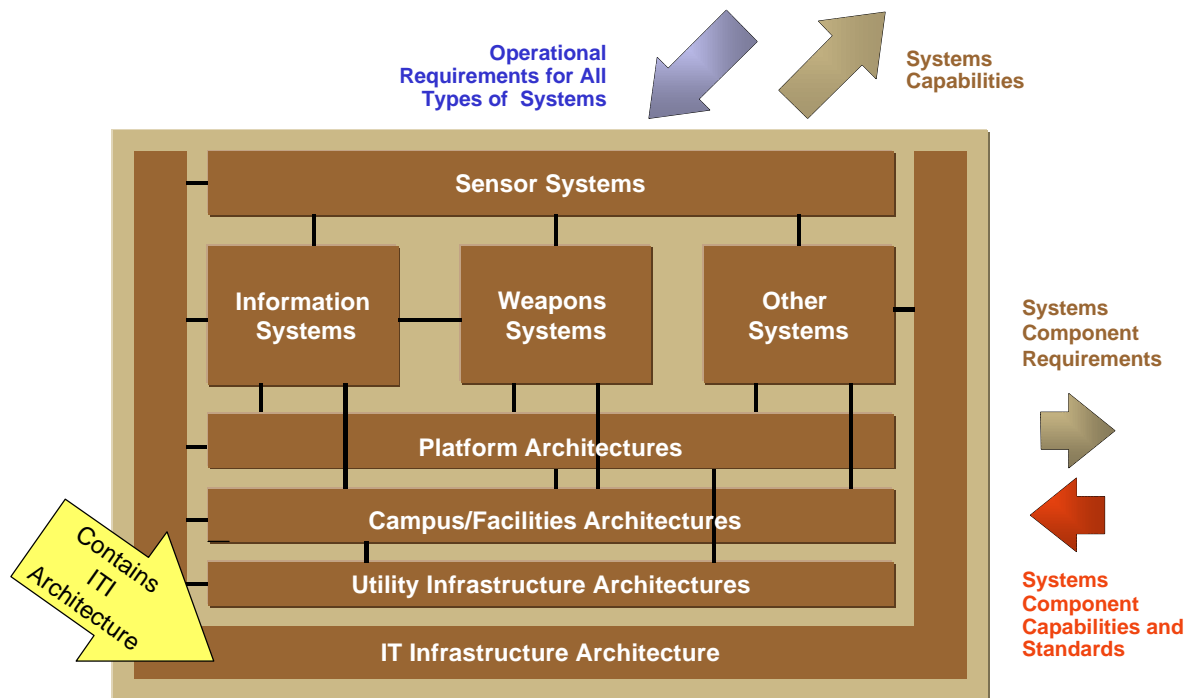


Figure 1-3. Systems View Showing Sub-Systems Views

There are a number of distinct but inter-related sub-systems which work together to produce the required operational capabilities.

- Sensor systems provide overall surveillance and data collection capabilities required to support the planners, commanders, and warfighters. Collectively, various sensor systems combine to produce information products required by various user communities.
- Information systems include all types of information processing and management applications. These are generally specific to Communities of Interest (e.g., logistics) or common across the DON.
- Weapon systems are a distinct class of systems including propulsion, guidance, and payload capabilities, all of which may imbed information technologies as part of their respective control systems.
- Other systems include simulators, trainers, robots, materiel handling systems, and other sensor-based and/or real time systems with special user interfaces.
- Platforms, including ships, planes, spacecraft, amphibious units, and vehicles, are also recognized as a distinct class of transportation systems. Similarly, campuses and facilities are also viewed as “land and building systems” with their own classes, components, and relationships. The utility infrastructure, including power, HVAC, water, and sewage, is further defined as a sub-view in support of platforms and campuses/facilities.

- The ITI Architecture supports all of the above sub-systems and directly interfaces with various user communities in the operating areas.

Figure 1-4 takes the bottom ITI Architecture Sub-view (of Figure 1-3) and drills down an additional level of detail. The components of the ITI are presented as the IT Infrastructure Sub-view. The ITI Architecture provides all of the common information access, management, and exchange services required by information systems and users of information technologies.

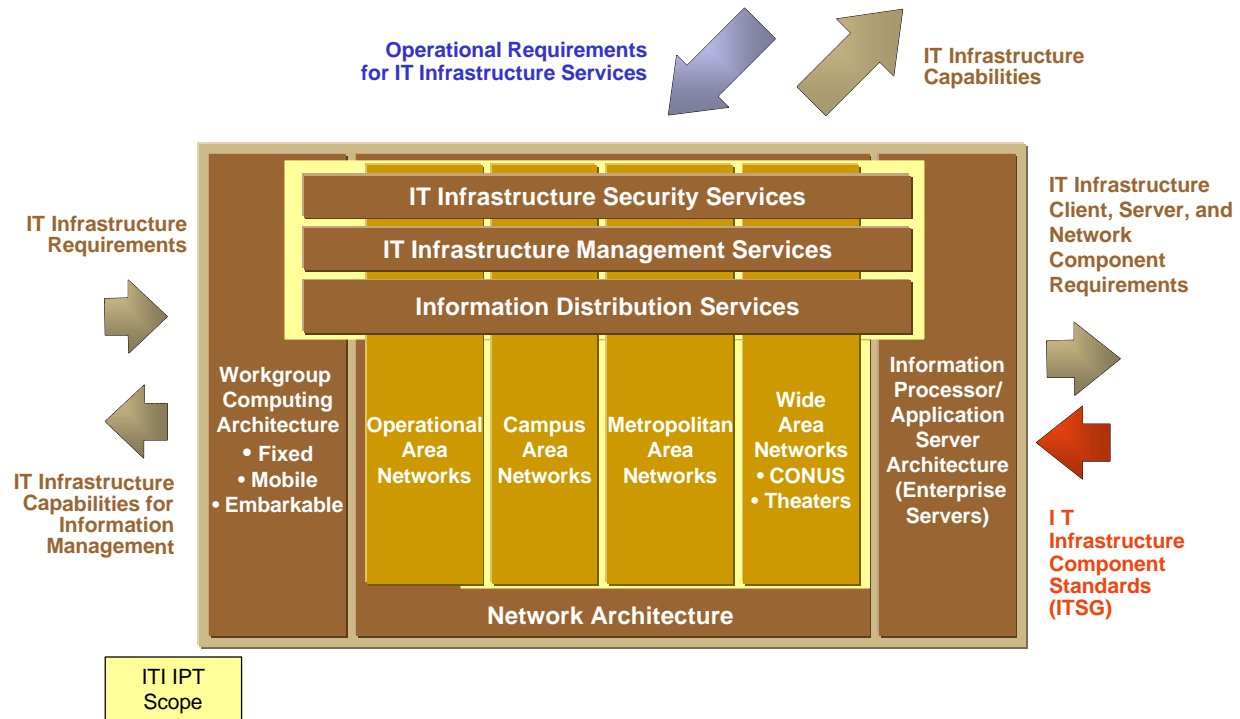


Figure 1-4. IT Infrastructure Sub-view

The three underlying component architectures are the Workgroup Computing Architecture, the Network Architecture, and Server Architecture. The Workgroup Computing Architecture addresses workstations, peripheral devices, workgroup servers, and LANs. The Server Architecture addresses the multi-tiered structure of information and application processors and storage devices. The Network Architecture addresses the connectivity between workgroup devices and servers across the three levels of networks supporting ashore, afloat, and expeditionary communities.

This underlying Network Architecture capability is overlaid with three end-to-end infrastructure services, including information distribution services, IT infrastructure management services, and IT infrastructure security services.

- Information distribution services provide for various types of information exchange and communication services across the ITI.
- IT infrastructure management services provide performance- and service-level management capabilities plus other operational services.
- IT infrastructure security services provide for the stringent requirements for controlled access and information protection across the ITI for the various levels of security.

1.4.2 From Framework to ITI Architecture

Volume I addresses the network portion of the ITI architecture highlighted in yellow in Figure 1-4. It includes ITI security services, ITI management services, and information distribution services in conjunction with Campus, Metropolitan, and Wide Area Networks. Operational Area Networks support platform and expeditionary networking requirements. Operational Area Networks are not fully addressed in this document. They will be the subject of a future ITI document version. The ITI Network Architecture deliverables are depicted in Figure 1-5.

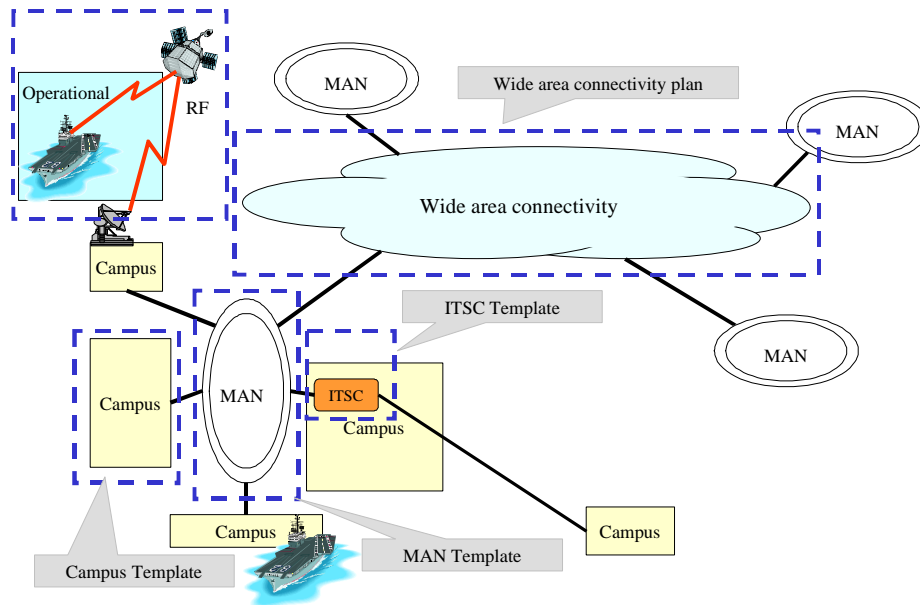


Figure 1-5 ITI. Architecture Templates and Plans

The deliverables include the following: Wide Area Network (WAN) Plan, Metropolitan Area Networks (MANs) Template, Campus Area Networks (CANs) Template, and Information Technology Service Centers (ITSCs) Template. These deliverables represent the building blocks of the enterprise network. These are addressed in detail the remaining chapters and appendices of this document.

This page intentionally left blank.

Volume I, Chapter 2 – Table of Contents

| | | |
|-------|--|-----|
| 2. | Network Infrastructure Requirements | 2-1 |
| 2.1 | Operational Requirements..... | 2-1 |
| 2.1.1 | All Communities of Users..... | 2-1 |
| 2.1.2 | All Operating Areas..... | 2-2 |
| 2.1.3 | All Types of Media..... | 2-2 |
| 2.1.4 | All Types of Application Networking Requirements | 2-2 |
| 2.1.5 | Customized Services to Meet Customer Needs..... | 2-3 |
| 2.1.6 | Operational and Management Support Capabilities | 2-4 |
| 2.1.7 | Evolution to Accommodate New Technologies | 2-4 |
| 2.1.8 | Cost Effectiveness | 2-4 |
| 2.2 | IT Network Infrastructure Functional Requirements..... | 2-5 |
| 2.2.1 | IT Infrastructure Security - Functional Requirements | 2-5 |
| 2.2.2 | IT Infrastructure Management - Functional Requirements | 2-6 |
| 2.2.3 | Information Distribution Service - Functional Requirements | 2-6 |
| 2.3 | IT Network Infrastructure Connectivity Requirements..... | 2-7 |
| 2.3.1 | Operational Area Network Connectivity Requirements | 2-7 |
| 2.3.2 | Campus Area Network Connectivity Requirements | 2-8 |
| 2.3.3 | Metropolitan Area Network Connectivity Requirements..... | 2-8 |
| 2.3.4 | Wide Area Network Connectivity Requirements | 2-8 |

This page intentionally left blank.

2. Network Infrastructure Requirements

The ITI must provide the underlying information transport and service networks for addressing the diverse operational requirements of the DON. This can be summarized as the capability to deliver accurate and timely information in the appropriate form to the intended recipients in a secure and cost-effective manner.

The ITI provides communication and information distribution capabilities to the DON and to its extended external user communities and locations. The magnitude of these required capabilities is difficult to accurately predict, given the rapidly changing demands and opportunities for leveraging IT for enterprise transformation. Therefore, a general requirement of the IT network infrastructure architecture is that it be flexible and scaleable in order to meet the expected rapid but unpredictable growth in all aspects.

The ITI for the DON must also support a wide variety of specific user requirements. The DON Board of Representatives established the requirement that the enterprise network architecture address tactical and non-tactical, afloat and ashore, Navy and Marine Corps applications. There are significant challenges (geographical, organizational, mobility, security, operational, budgetary, and technical) that must be specifically addressed and accommodated. Perhaps the most challenging is to support the myriad autonomous networks and communities of interest that currently operate within the DON and begin merging these into an integrated enterprise network capability.

The following is a list of DON ITI requirements organized according to the Network Architecture Framework. These requirements have been collected into a single chapter to provide a consolidated reference for the architectural solutions that follow in the remainder of the document.

2.1 Operational Requirements

The following lists represent general requirements of the operational capabilities of the DON IT network infrastructure.

2.1.1 All Communities of Users

The DON IT network infrastructure must meet the communication and connectivity needs of all parties involved in fulfilling the mission of the DON, including the following user communities:

- All commands, functions, and task groups within the DON
- Joint forces
- Allied forces
- Reserves
- Contractors
- Intelligence community
- Recruits
- Other government agencies
- Foreign governments
- Academic/research community
- Medical community
- General public
- Other stakeholders

2.1.2 All Operating Areas

The DON IT network infrastructure must meet these communication and connectivity needs in all operational areas and environments in an efficient and secure manner, including:

- CONUS bases and facilities of all types
- OCONUS bases and facilities of all types
- Battle Groups and ships at sea and in port
- Expeditionary forces at sea and on the ground
- Squadrons and planes on base, on ship, and in the air
- Any facilities, equipment, or personnel deployed in space
- Remote access for mobile users
- Access to and from contractor sites
- Access to and from telecommuting and home locations
- CONUS and globally-distributed functions and teams

2.1.3 All Types of Media

With ongoing advances in digital networking, the DON IT network infrastructure is expected to support all types and forms of information transport in secure and non-secure transmissions, including:

- Data , including numeric, text, graphics, and images
- Voice
- Video
- Multi-media
- Non-POTS telephony services (mobile, cellular, and radio frequency), digital telephony, voice mail, computer-integrated telephony, interactive voice response, single-line concept, and enhanced 911
- Paging

2.1.4 All Types of Application Networking Requirements

There are many different types of systems that do or will depend upon the DON IT network infrastructure for meeting their communication requirements. There are networking requirements for data gathering, online access for users, inter-application information exchange, and the use of the network for distributed database access. These networking requirements for applications include:

- Information system application communications for all communities
- Support for the Global Command and Control System (GCCS) and the Global Combat Support System (GCSS)
- Support for the common tactical picture

- Support for damage control information (both conventionally, i.e. detecting damage inside a ship, and in a wider sense, i.e. detecting a chemical plume to establish decontamination boundaries)
- Support for sensor-to-shooter (especially important for Marines with a decreased footprint ashore who rely increasingly on offshore shooting assets for artillery support)
- Support for distributed combat systems such as TBM Defense
- Support for facility monitoring and control systems
- Network capacity to support data warehousing and data mining
- Supports quality of life for our sailors and Marines
- Support for approved legacy protocols consistent with ITSG
- Support for migration from legacy to adopted protocols
- Support for access to information/data maintained at and provided by non-government sites
- Phased support for force coordination, force control and critical support, and weapons control and systems control and monitoring
- Guidance and support for application planners and developers on designing for effective use of network capabilities
- Supports Y2K compliance

2.1.5 Customized Services to Meet Customer Needs

Various organizations and user communities have the requirement to establish autonomous or virtual networks within the DON IT network infrastructure. These autonomous or virtual networks must support the organizations' and communities' special needs for security, performance, and functionality.

The planning, design, and provisioning of enterprise networking services must provide the flexibility to customize services to meet specific customer needs. In order to leverage economies of scale and provide a standardized and dependable global infrastructure for all user communities, it is recommended that a common set of basic services be provided across the enterprise. The definition of basic and optional services is provided below.

- **Basic Services:** Those services defined as containing the aspects of a utility, including essential need, benefit from sharing, generic in nature, specificity not essential for individual requirements, subject to economies of scale, and can be reliably provided to many. Also includes services that should be regulated for some reason, including price, quality of service, safety, security, required investment, and law or policy.
- **Optional Services:** Those services that contain the aspects of uniqueness, individuality, or tailoring needed to meet specific needs and requirements of customers. These are services that a significant portion of the customer base does not regard as "basic" or "utility" services or capabilities.

The resulting package of network services and service levels will be established through the use of Customer Service Level Agreements. These agreements define the nature of associated use and performance measurements and will be linked to pricing and billing arrangements for each customer of the IT infrastructure.

2.1.6 Operational and Management Support Capabilities

Network customers and users require the following ongoing operational and management support capabilities from the Infrastructure Service Provider(s) which must be supported by the IT network infrastructure.

- Efficient and responsive service ordering and provisioning
- Troubleshooting and help support for installation, set-up, and on-going operations
- Timely identification and resolution of problems
- An integrated billing system (including phone support) based on services acquired and use of the infrastructure
- Detailed accounting for services provided
- Provision of information and the capability to collect measurements, including:
 - ♦ Service use, including peak or surge characteristics, and available capacity
 - ♦ Performance criteria (bandwidth, latency, Quality of Service)
 - ♦ Service availability and outages
 - ♦ Survivability and timely casualty restoration
 - ♦ Service performance measurements

2.1.7 Evolution to Accommodate New Technologies

The DON IT network infrastructure will be built using current technologies with acceptable and manageable technical and economic risks. The architecture and resulting choices of technology components will evolve to embrace new and emerging technologies as these are seen to offer discernible advantages to customers. The means to update and refresh the architecture guidance and resulting design specifications must be built into the operational model for ongoing planning and management of the DON IT network infrastructure. This requirement includes:

- Processes for collecting future user requirements and providing feedback to planners and developers
- Support of Advanced Concept Technology Demonstrations (ACTD) and Advanced Technology Demonstrations (ATD) involving users

2.1.8 Cost Effectiveness

Providing and operating the DON IT network infrastructure must be based on sound business case analysis and management practices to optimize the ability to meet user requirements with the associated costs. This implies:

- Application of best commercial practices
- Pricing (or costs) comparable to those in private industry
- Ability to choose service providers that can deliver the best value
- The use of the Total Cost of Ownership (TCO) approach for all cost evaluations

2.2 IT Network Infrastructure Functional Requirements

In addition to the above general requirements to support the systems and operational requirements of the DON, more specific functional requirements have been identified. These have been defined using the relevant components of the EAF for network architecture.

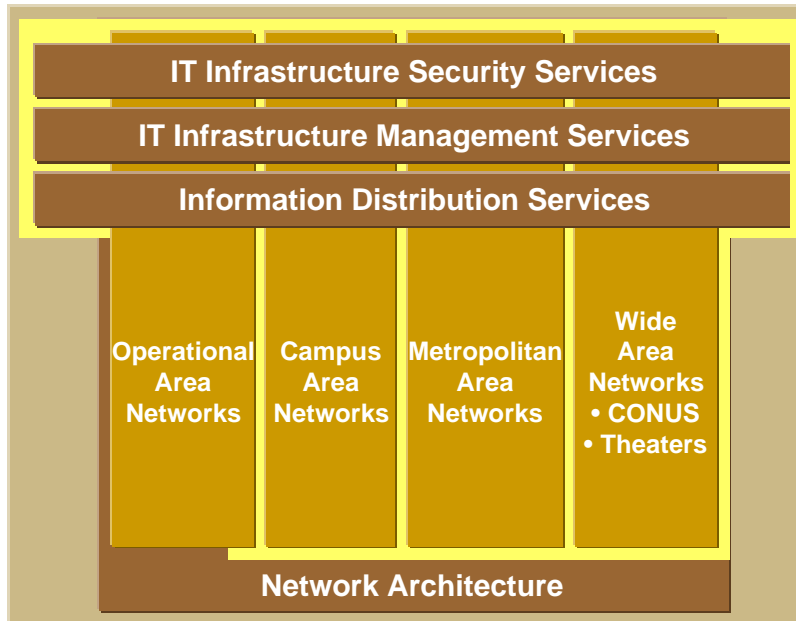


Figure 2-1. Framework for Identifying IT Network Infrastructure Requirements

The DON IT network infrastructure functional requirements are identified in the following sub-sections for each of the three service levels in this framework. Specific network connectivity requirements for each of the four classes of networks are further identified in the subsequent section.

2.2.1 IT Infrastructure Security - Functional Requirements

One of the highest priorities for DoD and DON global networks is security. There is a requirement for the DoD “Defense in Depth” program to be fully implemented in the DON. This will assure confidentiality, integrity, authentication, and non-repudiation. All security requirements must be met across all networks, including the attached clients and servers. The following detail the specific network security services that are required:

- Access control services
- Intrusion detection services
- Encryption services
- Attack shunning
- Public key administration
- Proxy services
- Vulnerability detection services
- Malicious code detection
- Illegal use detection
- Enclave protection

2.2.2 IT Infrastructure Management - Functional Requirements

There is a need for end-to-end global distributed network management. This includes the requirements for a series of integrated network operations centers. The IT infrastructure management requirements are:

- Fault management
- Performance management
- Configuration management
- Software distribution
- Help desk services
- Usage accounting management
- Capacity planning, network modeling, and trend analysis
- Asset management
- Integration management
- File management

2.2.3 Information Distribution Service - Functional Requirements

The Information Distribution Services provide IT infrastructure users with various capabilities to manage the distribution and delivery of information. These requirements are defined below. They are divided into basic and optional services.

Basic Services include:

- Network time service
- Domain name/network addressing services
- Enterprise directory services
 - ♦ Directory synchronization within the DON
 - ♦ Directory service accessibility by applications and systems
- Message transfer services
- Electronic mail (e-mail)
- E-mail attachments
- Network news service
- Web-hosting and transport services
- File transfer services
- Remote access services
- General voice (including conferencing)
- Shipboard voice
- Secure voice
- Multimedia services

Optional Distribution Services include:

- Defense Messaging System
- Facsimile transmission
- Electronic Commerce/Electronic Data Interchange (EC/EDI)
- Workgroup computing
- Electronic dialog (chat)
- Collaborative planning with imagery (groupware)
- Video Teleconferencing (VTC)
- Multi-cast broadcasting and other “many-to-many” and “one-to-many” applications
 - ♦ Software distribution services
- Television broadcasting
- Software distribution services

2.3 IT Network Infrastructure Connectivity Requirements

The DON IT network infrastructure must provide end-to-end connectivity across all four network categories – operational, campus, metropolitan, and wide area. This includes connectivity to all users and operational locations as outlined above, including internal and external users/locations and remote/mobile access.

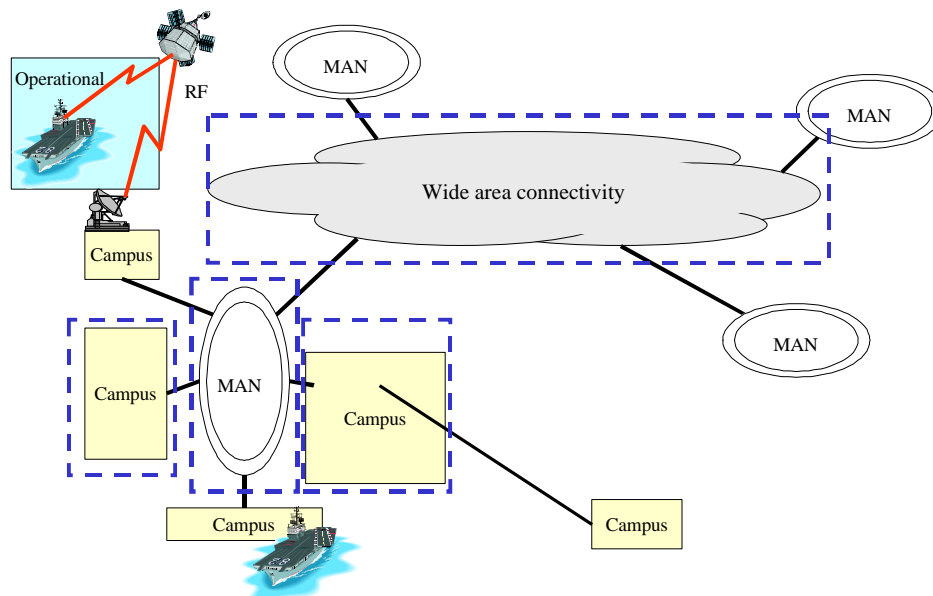


Figure 2-2. End-to-End Network Connectivity

The specific inter-network connectivity requirements are summarized in the following four sub-sections for each of the four network categories.

2.3.1 Operational Area Network Connectivity Requirements

There is a requirement to optimize operational area services to enable a solution to the fleet mobility problem and to enable the fleet intranet. The Operational Area Network connectivity requirements include:

- Military satellite communications
- Commercial satellite communications
- Communication link system
- COMSECURITY system
- Military line-of-sight communications
- Commercial line-of-sight systems
- Communication transceivers
- Communication Umbilical System
- Fleet Teleport System
- Communication Antenna System
- Pier interface

2.3.2 Campus Area Network Connectivity Requirements

There is a requirement within the campus area network to optimize campus area services. The CAN connectivity requirements are:

- Campus Area Network
- Campus Area Network/Local Area Network interfaces
- Campus Area Network/Metropolitan Area Network interfaces
- Campus Area Network/Wide Area Network interfaces

2.3.3 Metropolitan Area Network Connectivity Requirements

There is a requirement within the metropolitan area network to optimize the provisioning of metropolitan area services in areas of high concentration of DON and joint campuses. The MAN requirements are:

- Metropolitan Area Network
- Metropolitan Area Network/Local Area Network interfaces
- Metropolitan Area Network/Campus Area Network interfaces
- Metropolitan Area Network/Wide Area Network interfaces

2.3.4 Wide Area Network Connectivity Requirements

There is a requirement within the wide area network to optimize wide area services. The WAN requirements are:

- Wide Area Networks
- DISN
 - NIPRNET
 - SIPRNET
- Internet
- Commercial telephone
- Secure commercial telephone
- Defense Switched Network (DSN) access
- Wide Area Network/Local Area Network interfaces
- Wide Area Network/Campus Area Network interfaces
- Wide Area Network/Metropolitan Area Network interfaces
- Wide Area Network to external network interfaces

Volume I, Chapter 3 – Table of Contents

| | |
|---|------------|
| 3. Conceptual ITI Network Architecture..... | 3-1 |
| 3.1 ITI Network Architecture Approach..... | 3-1 |
| 3.1.1 Geographic Considerations | 3-2 |
| 3.1.2 Connectivity Grid | 3-2 |
| 3.1.3 Autonomous Networks | 3-3 |
| 3.1.4 Communities of Interest (CoI)..... | 3-3 |
| 3.1.5 Users | 3-3 |
| 3.2 Layered Technical Model | 3-4 |
| 3.3 ATM Connectivity Architecture Overview | 3-5 |
| 3.3.1 DON ATM Architecture Strategy..... | 3-5 |
| 3.3.2 ATM Connectivity Approaches..... | 3-6 |
| 3.3.3 ATM Addressing Plan | 3-8 |
| 3.3.4 ATM Routing Architecture | 3-9 |
| 3.3.5 Rationale for ATM Technology in DON ITI Architecture..... | 3-9 |
| 3.4 IP Connectivity Architecture Overview | 3-11 |
| 3.4.1 Strategy | 3-11 |
| 3.4.2 IP Connectivity | 3-11 |
| 3.4.3 Autonomous Networks | 3-12 |
| 3.4.4 Fleet Intranet | 3-13 |
| 3.4.5 IP Addressing Plan | 3-13 |
| 3.4.6 IP Routing Architecture | 3-14 |
| 3.5 Network Connectivity Security Overview..... | 3-15 |
| 3.6 ATM Connectivity - Detailed Architecture..... | 3-17 |
| 3.6.1 ATM Planning and Implementation Constraints | 3-17 |
| 3.6.2 ATM Architecture Design Factors..... | 3-18 |
| 3.6.3 ATM Addressing and Routing | 3-18 |
| 3.6.4 ATM Protocols | 3-21 |
| 3.6.5 ATM Overlay Security..... | 3-25 |
| 3.7 IP Connectivity Detailed Architecture | 3-27 |
| 3.7.1 IP Connectivity Design Factors | 3-28 |
| 3.7.2 Placement of Routers in IP Connectivity | 3-28 |
| 3.7.3 IP Addressing | 3-30 |
| 3.7.4 IP Routing | 3-30 |
| 3.7.5 IP Overlay Security..... | 3-31 |
| 3.7.6 Considerations for Connecting Contractors..... | 3-32 |
| 3.8 DON ITI Architecture Plan of Action..... | 3-33 |

Draft Working Papers of the ITI IPT

| | | |
|-------|--|------|
| 3.8.1 | Steps for Developing Detailed Enterprise ITI Architecture..... | 3-33 |
| 3.8.2 | Steps for Developing a MAN | 3-35 |
| 3.8.3 | Steps for Developing a CAN..... | 3-36 |

3. Conceptual ITI Network Architecture

Network connectivity is a fundamental requirement for a DON integrated enterprise information infrastructure. This chapter defines the ITI architecture that will support planning and implementation of the network connectivity segment of the infrastructure. The ITI planners address connectivity requirements by using architecture plans and templates that guide and constrain their solutions. The result is a set of network connectivity solutions that are consistent, complementary, and interoperable and that support the connectivity requirements of the DON functional missions.

3.1 ITI Network Architecture Approach

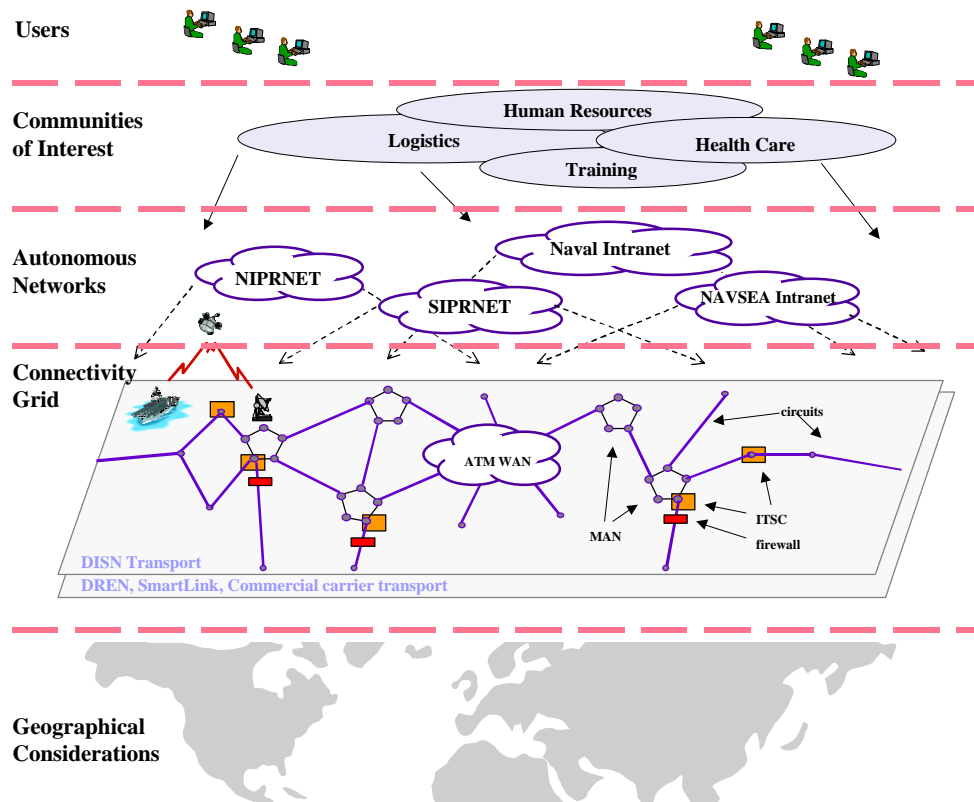


Figure 3-1. Conceptual Network Architecture

The development of the Naval ITI Network Architecture was predicated on a number of concepts, considerations, and doctrines associated with the Navy and Marine Corps mission, organizations, and operating relationships. Figure 3-1 is a composite of these factors and provides a simple view of the approach taken to develop this architecture. The five layers of the Conceptual ITI Network Architecture and the interrelationships of the layers are described in the following sections.

3.1.1 Geographic Considerations

Most of the Naval activities and personnel are concentrated at specific locations around the globe. These areas of concentration are illustrated in Figure 3-2.

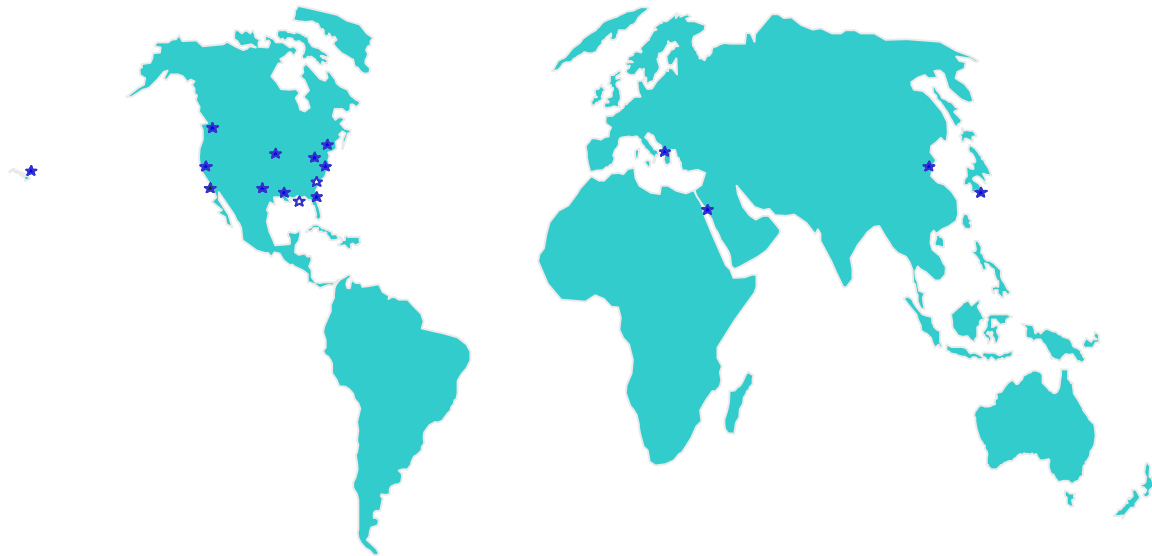


Figure 3-2. Naval Concentration Areas

These areas of concentration tend to align with metropolitan areas. Metropolitan areas typically have a high population density and an advanced telecommunications infrastructure. It is also predictable that they will have high bandwidth available and that it will be relatively inexpensive. A natural progression is to implement MANs in these areas to provide rich connectivity between Naval activities within the region.

Afloat units present a very different set of geographic considerations because of their dispersed, highly mobile, and wide-ranging deployment patterns. The satellite communications downlinks in support of afloat units are at well-defined locations around the world. When in port, shipboard units become aligned with metropolitan areas.

3.1.2 Connectivity Grid

In Figure 3-1, there is an underlying connectivity grid or common transmission fabric (includes the circuits, area networks, switches and routers, and RF links) that provides raw bandwidth, and on top of which are built all network services. The MANs are interconnected with wide area connectivity (networks or circuits). Campus networks outside the metropolitan area are linked via point-to-point links or other means. Each region has an Information Technology Support Center (ITSC) connected via the MAN. Fleet teleports are connected to the enterprise network in the same way that a base is connected (on the MAN or via a circuit). Piers are connected in a similar way, but are generally part of a base and are connected via the associated campus network. Shipboard connectivity is provided via an RF communications infrastructure when afloat or via the pier when ashore.

The communications infrastructure is built primarily of components that are external to the Navy. In particular, the WAN and MAN connectivity is provided through a combination of telecommunications carriers and DoD- and/or Naval-owned devices. The campus connectivity, on the other hand, is primarily owned by the Navy and Marine Corps.

There are other global WANs to which the DON enterprise network must interconnect. Examples include NIPRnet, SIPRnet, and the Internet.

3.1.3 Autonomous Networks

Autonomous networks are built on top of the connectivity grid. Autonomous networks are generally independent of each other from the standpoint of media, technology, security, management, and other characteristics. They are more than virtual networks and include physical components that are unique to the particular autonomous network.

Autonomous network examples include the fleet intranet, the Naval Intranet, the SYSCOM networks, classified overlay networks, and voice and video networks.

3.1.4 Communities of Interest (CoI)

This layer represents a user-centric, geographically-dispersed grouping at or above the autonomous network layer. An example of a CoI is any functional area, such as logistics, where users need to exchange information relating to that functional area, and this exchange is across the user's normal organizational boundaries. The CoI, in cases such as logistics, will have extensive associated information management and systems applications that all DON users access to work in the functional area. A given CoI must have access to multiple autonomous networks to obtain the required connectivity services. This CoI capability forms a foundation requirement for enabling RBA and RMA.

Other CoI examples include intelligence, human resources, acquisition, training, and health care. Some, if not many, of these users may be members of multiple communities of interest.

3.1.5 Users

The users of the DON enterprise network include all Naval military and civilian personnel. It also includes contractors and other support personnel, as well as other parties that have some association with Naval activities. Because users are typically in multiple CoIs on physically diverse autonomous networks (as described above), Virtual Private Networks (VPNs) must be supported by the infrastructure.

3.2 Layered Technical Model

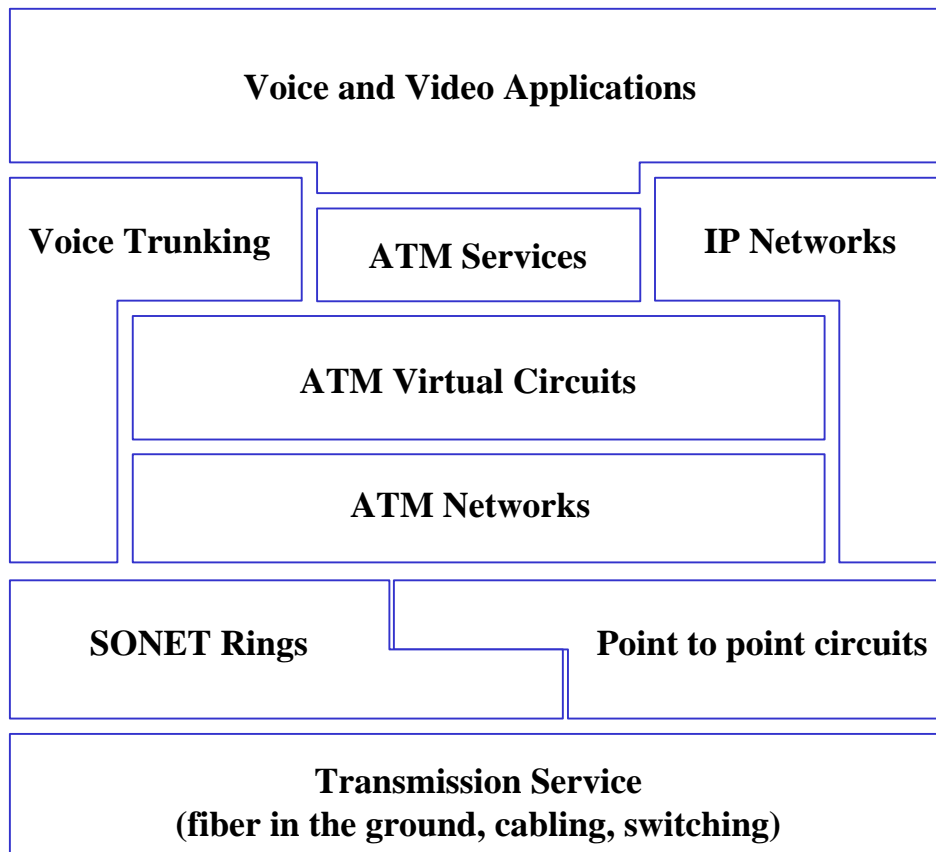


Figure 3-3. A Layered Approach to Networking Solutions

A layered technical model, depicted in Figure 3-3, was adopted as a basis for considering, organizing, developing, and presenting the ITI architecture. The IP network-centric view starts at the bottom with the transmission fabric (fiber, copper, multiplexers, etc.); on top of which are SONET rings or point-to-point circuits; on top of which are ATM networks; on top of which are ATM Virtual Circuits or links using Switched Virtual Circuits (SVCs), Permanent Virtual Paths (PVPs), or Permanent Virtual Circuits (PVCs); on top of which are IP networks. Some of the layers are optional, depending on the requirements of the autonomous networks or the communities of interest that use them.

3.3 ATM Connectivity Architecture Overview

This section presents the conceptual ATM connectivity architecture. Greater detail is provided for planners and implementers in Section 3.6. It is assumed that the reader has basic familiarity with ATM terminology, standards, and protocols.

3.3.1 DON ATM Architecture Strategy



The ATM architecture must support end-to-end ATM cell delivery and signaling. End-to-end is defined as from campus-to-campus or, in some cases, a campus may deliver ATM service deeper into the local infrastructure (in some cases to the desktop). In the latter cases, end-to-end service means supporting desktop-to-desktop ATM connectivity. To clarify, end-to-end service means that end-to-end signaling is supported to provide SVC service. In a large, complex network, this connectivity and service must be supported by a well-defined addressing and routing plan.

In the current ATM service provider market, each service provider would prefer that the customer use the service provider's addressing and routing architecture. This is acceptable except when customers are dual-homed to different service providers, such as in geographically-dispersed organizations. As a result, the customers will be compatible with one service provider but not others. This is particularly important in the DON because the entire DON enterprise network is comprised of many ATM service providers at both the MAN and WAN levels. It is a certainty that the DON will be inconsistent with some service provider's routing and addressing plans, no matter which plan is chosen.

The ATM architecture described here is independent of any specific ATM service provider and provides a mechanism to route and signal through any ATM network. The mechanism to accomplish this is "tunneling" through the various MAN and WAN ATM networks, which requires use of PVPs to create VPNs over the ATM service provider networks.

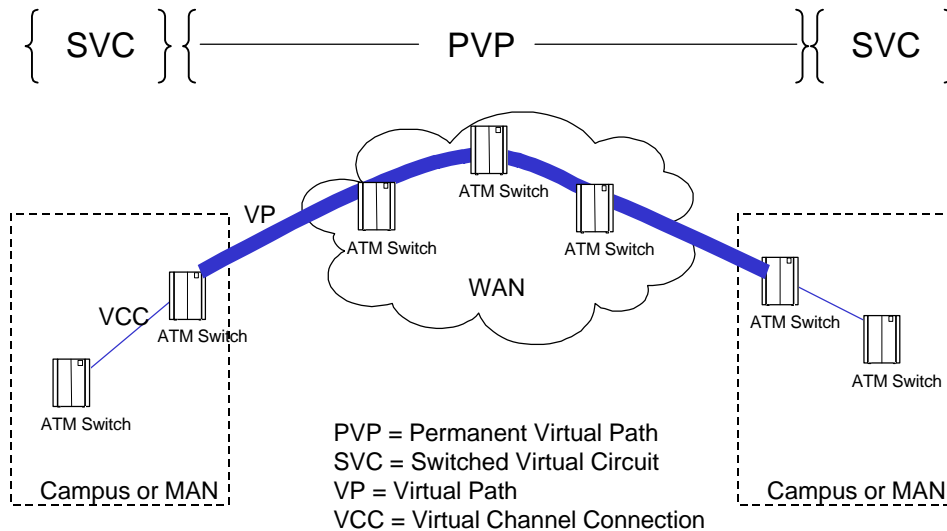


Figure 3-4. Combining PVPs and SVCs in an Architecture Solution

The VPN solution is depicted in Figure 3-4 and specifies a PVP mesh. It is for use at the levels of the network hierarchy in which the DON addressing and routing architecture are incompatible with that of the ATM service provider. The PVP mesh provides a set of fixed conduits through which the DON can establish a signaling and routing environment that supports SVCs and also support provisioning of PVCs. The advantages here are that Naval managers can still obtain SVC service from the PVP end points and they can easily add or modify PVCs as requirements change. With this approach, changes to the PVP mesh, and eventually the signaled environment, should have the least impact on campus and end users. The following PVP guidance applies.

- **WANs.** If the DON uses the DISA ATM addressing plan, and if the Wide Area Connectivity solution is the DISN ATM Service, then a PVP mesh will not be required for the general WAN connectivity case because the DON will be compatible with the provider's addressing conventions. However, in order to support a highly-mobile fleet autonomous ATM network, a PVP mesh will need to be established to support the fleet requirements.
- **MANs.** At the MAN level, some type of PVP mesh will be required as an overlay to any commercial ATM service offered as part of the MAN.
- **CANs.** At the campus level, the PVP solution should not be used. The campus networks are Naval owned and operated and therefore will comply with DON addressing and routing standards.

3.3.2 ATM Connectivity Approaches

There are two basic ATM connectivity cases that must be accommodated by this architecture. One is the general case of connecting the campus networks with the MANs with further connectivity to the WAN. This is a general (non-mobile) geographic situation. The other is the fleet case, which requires a high degree of mobility and must be supported by a dynamic routing architecture. The architecture presented here fully supports both of these environments. (The

intermediate case of “portable” users, in which service is not maintained during transit, is covered within.)

3.3.2.1 General (Non-Mobile) Case

The ATM switch that links the campus infrastructure to the external world (MAN) is referred to in the DON ITI architecture as a “premise” switch (or Primary Information Transfer Node (PITN) in BLII terminology).

The campus premise switch will connect to the MAN, preferably at the 155-Mbps OC-3c level of bandwidth (Figure 3-5 refers). Independent of the selected MAN technology (point-to-point, SONET, or ATM), the requirements for signaling and routing will drive the MAN architecture to a DON-controlled and -managed routing and signaling domain. For an ATM MAN, a PVP mesh overlay will be implemented to support the signaling domain between campuses.

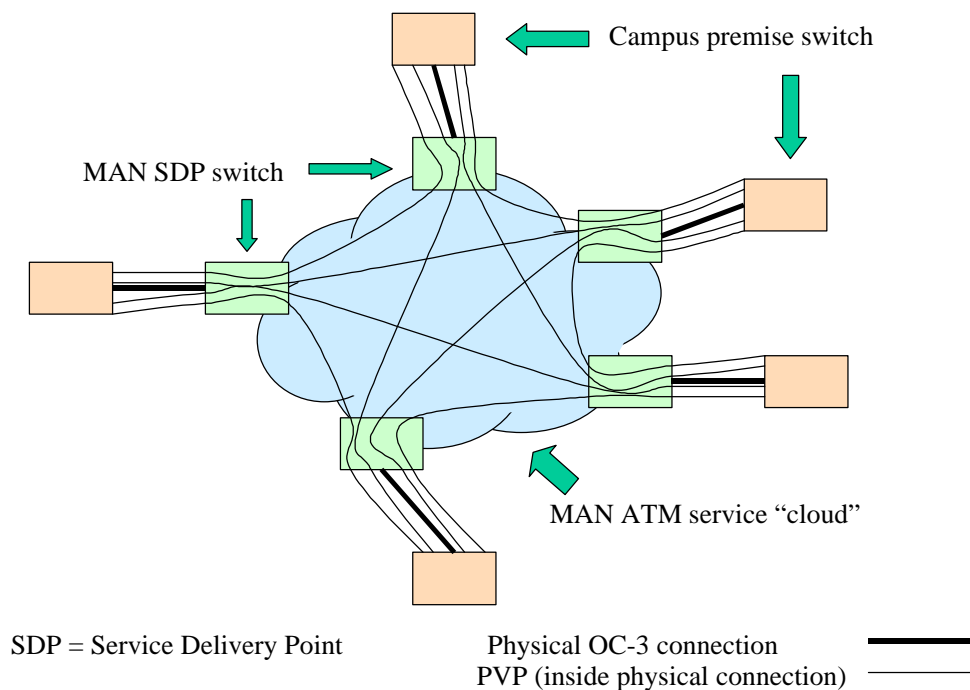


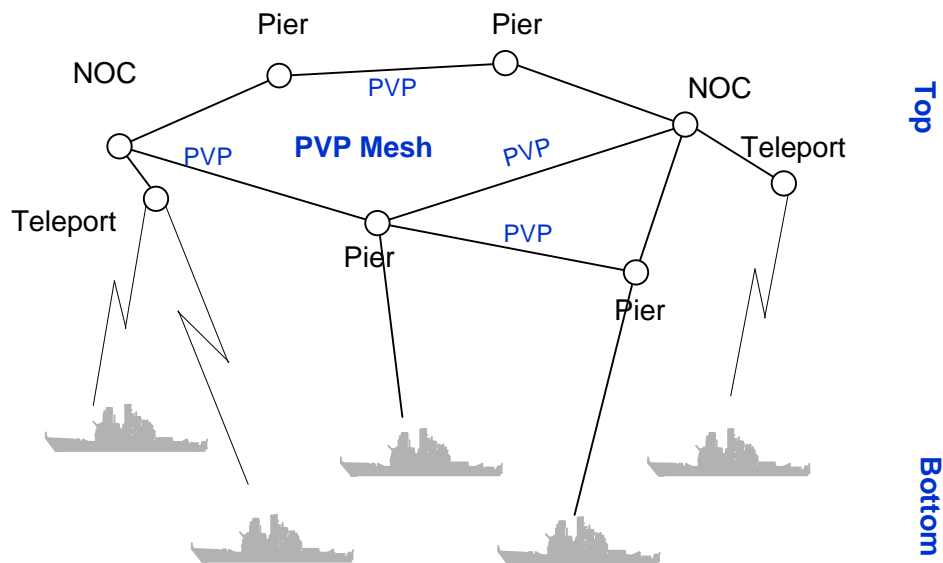
Figure 3-5. MAN Service delivered as ATM, with PVP Mesh to Campus Premise Switches

The ITSC premise switch normally will also have connectivity to the ATM WAN in order to establish connectivity with other regions. There are exceptions to this rule. For example, if the ITSC is at one campus and the WAN provider already has a POP at another campus, there may be no reason to install a new WAN POP at the ITSC if there is appropriate MAN access at the other campus premise switch.

3.3.2.2 Fleet (Mobile) Case

This section only applies to ATM-enabled ships that wish to have a native ATM connection to the rest of the fleet intranet. It does not apply to ships that have no external ATM connections - that is, those ships that are IP only.

In the fleet ATM connectivity case, our Naval mission requires full support for mobility. While fully-signaled mobile connections are desired, the standards (and existing implementations) are not yet developed. Until such time, there will be a separate fleet autonomous network based on ATM technology, with a PVP mesh linking the piers, Standard Tactical Entry Point (STEP) sites, and fleet NOCs. The mesh provides end-to-end signaling and cell delivery and allows easy reconfiguration as the mobile network moves. This does not require a full mesh because the supporting PVPs follow actual physical paths. As shown in Figure 3-6, routing is two-level PNNI with the ashore PNNI mesh and fleet switches at the top level, and the afloat units at the bottom. The boundary between the levels is at the RF links and the pier connections. This signaling domain does not extend to other Naval or even to DoD signaling domains. The IP connectivity must be separated by firewalls anyway. If signaling outside these constraints is required, it will be engineered on a case-by-case basis until a general solution becomes available.



PVP = Permanent Virtual Path

NOC = Network Operating Center

Figure 3-6. Two-level PNNI Hierarchy

Because the mobile platforms house other deployable forces such as MEUs, provisions for this multi-layer deployment must be addressed.

3.3.3 ATM Addressing Plan

The DON Connectivity Architecture incorporates a geographically-based addressing plan. To support DoD conventions and allow long-term interoperability within DoD, DON will use the DISA ATM addressing plan whenever possible. For mobile users, DISA has offered Globally Unique Identifier (GUI) addresses that will follow the deployed forces. The GUI addresses do not meet the DON requirements and will not be used.

A MAN will obtain an ATM address block to support the MAN and all campuses within the region. Implementation will deviate slightly from the DISA address plan in that a single (or small) number of DISN nodes will be assigned in a Naval MAN. Addressing behind this node(s) will be performed at the discretion of the MAN architect. The hierarchy of the DISA plan will be maintained outside the MAN. The MAN will allocate address space from the DISA block to the campuses.

Fleet addresses will be allocated to ships based on their home port. A Naval base that includes piers should obtain enough campus network address space to allocate address prefixes to each ship that is homeported at that campus location.

Additional information on the DON ATM Addressing Plan is provided in Section 3.6.3.1.

3.3.4 ATM Routing Architecture

To exchange topology information for routing purposes, PNNI will be used whenever possible. Three levels of routing hierarchy are envisioned – campus, regional (MAN), and global. The manner in which this is achieved is dependent on the providers and technologies used at each of these levels.

From the MAN perspective, the MAN routing architecture must support routing between campuses. There will be limited prefix/masks per campus, but there may be multiple links between the campus and the MAN (for redundancy). PNNI will be the MAN routing protocol, and each campus will be viewed as a “logical peer group node.” The campus premise switch will “advertise” the appropriate campus prefix/mask to the MAN. When the campus ATM network has a rich topology, it may want to participate in the PNNI routing domain. In this case, it will be lower in the PNNI hierarchy. The choice to participate in the PNNI routing domain or to have prefixes/masks configured statically is determined by the MAN region/campus manager. In all cases, each campus will have a default route to the MAN, and each MAN will have a default route to the WAN.

Operation of the WAN requires that the prefix/mask information be known for each of the MANs and for any other entities to which it connects. This information can be inserted statically, or it can be derived through the normal operation of the PNNI protocol. Some type of dynamic scheme may be required because MANs need to have redundant paths to the WAN and only a dynamic routing protocol will allow re-routing in case of a connection failure.

Additional information on the ATM Routing Architecture is provided in Section 3.6.3.2.

3.3.5 Rationale for ATM Technology in DON ITI Architecture

Alternatives to ATM were considered during the selection of the networking technology; the ITI IPT selected ATM technology based upon a number of factors.

Today, time division multiplexing (TDM) technology is used in many Naval applications (using multiplexers) to provide multiple virtual connections over a single physical channel. Bandwidth allocation is fixed within each virtual channel, regardless of whether it is carrying any data, so it is impossible to “burst” to higher bandwidths on demand. ATM technology, on the other hand, provides the capability of multiplexing virtual channels and offers the efficiency of statistical

multiplexing in which bandwidth is consumed only when needed and is shared on a demand basis across the entire physical channel. This allows bursting to the physical line rate or to whatever bandwidth is unused by other channels. As a result, ATM provides significant efficiency gains over TDM technology, even with ATM technology's 10 percent overhead for header information.

DON requires that:

- ATM connectivity exists all the way to the desktop in some Naval network installations that have implemented IT21. This requires end-to-end ATM connectivity, even over the wide area.
- Multiple autonomous or virtual networks operate over the DON physical network infrastructure. ATM technology offers support for virtual channels with various classes of service and this provides the enabling capability for constructing the virtual networks. Alternatively, IP technology can be used to construct virtual networks on top of IP networks, but the IP approach is less efficient (it may increase overhead up to 50 percent) and fails to offer some required classes of service, such as circuit emulation.
- Networks provide a constant bit rate (CBR) or circuit emulation service for carrying synchronous trunks over a common communications fabric. Applications must include voice trunking between PBXs and bulk encrypted links using traditional type 1 encryption devices such as KG-84. ATM provides this class of service.
- Fast end-to-end key-agile (as opposed to bulk) type 1 encryption carry classified traffic over the common network. The one device that can achieve that operates over ATM networks (the KG-75).
- Significant scalability exists. ATM provides this without the need to replace the transport media (fiber).

The ATM decision is also examined in terms of other competing technologies. Compared to ATM, fast Ethernet and gigabit Ethernet technologies are better understood, appear to be scaleable, and are less expensive per drop. However, ATM was chosen for the "core" network technology, not necessarily for the "edge" networks. The appropriate technology solutions for connecting edge systems (e.g., hubs) are a separate issue. The advantages of ATM in the "core" network are exhibited on the backbone (wide area, metropolitan area, and some campuses) where scaleability is needed and many communications channels can be serviced over a single physical network. Additional advantages of ATM over the other technologies are addressed in the following.

- Ethernet Frame Size. The Ethernet class of protocols has a maximum frame size of only 1500 bytes. This frame size is too small to work well on networks near the border of the wide area and on the WAN itself. If the technology on or near the WAN allows a larger frame size, then fragmentation of the packets will occur at the edge networks, and this will cause significant performance degradation on the WAN. The DON architecture must eliminate such problems because there is no way to engineer around these limitations. Hence, we have the goal of delivering ATM service to each campus from the wide area or metropolitan area networks.

- Ethernet Scaleability. With ATM, only the optics at the interface level must be upgraded to achieve a faster optical carrier (OC) rate. With Ethernet, different cabling may be needed to migrate to fast Ethernet, FDDI, or even gigabit Ethernet if the current copper cable plant is inadequate for the higher bandwidths.
- Packet over SONET (POS). POS or Packet over wave division multiplexing (WDM) have been the technologies of choice in a number of new high speed demonstration networks that are not using ATM. While these new transport protocols offer significant promise for the future, they currently lack the maturity as well as other features and robustness required to meet Naval requirements.

3.4 IP Connectivity Architecture Overview

The conceptual IP connectivity architecture is presented here. Greater detail useful for planners and implementers is described later in Section 3.7. It is assumed that the reader is familiar with IP terminology, standards, and protocols.

3.4.1 Strategy

This IP architecture must support ubiquitous, end-to-end IP connectivity within the enterprise network and must support full access to the global Internet. Determining the IP architecture is relatively straightforward because the DON has significant experience upon which to draw. Moreover, the conventions, protocols, and implementations are mature.

The IP challenge is to provide an architecture that accommodates the separate autonomous networks, particularly the “fleet intranet”. The underlying ATM infrastructure enables the construction of these autonomous networks using virtual circuits as links between the routers in a given autonomous network. Each of these autonomous networks is a separate routing domain unto itself, and the capability must exist to route IP traffic from one autonomous network to another and to the external world. An additional constraint is that some autonomous networks must be protected from other autonomous networks using firewalls or similar technology. The fleet intranet introduces the necessity and challenge of supporting network mobility. Ships constantly change geographic positions, and their connectivity is dynamic both in topology and in bandwidth. Ships also require a strong security perimeter. For these reasons, there will be a specific architecture discussion for the mobile networks.

In general, there will be a minimum of detail regarding the various autonomous networks. The emphasis is on enabling the implementation but not to dictate the internal architecture of each network. That is left up to those who assess the requirements and develop the specific design for the individual autonomous networks.

3.4.2 IP Connectivity

3.4.2.1 General Case

The campus router that connects to the MAN will be termed the “premise” router for the purposes of this discussion. Figure 3-7 shows the relationships of the premise router connecting an

autonomous network to the MAN. In actuality, this router physically connects to the premise ATM switch over a direct link.

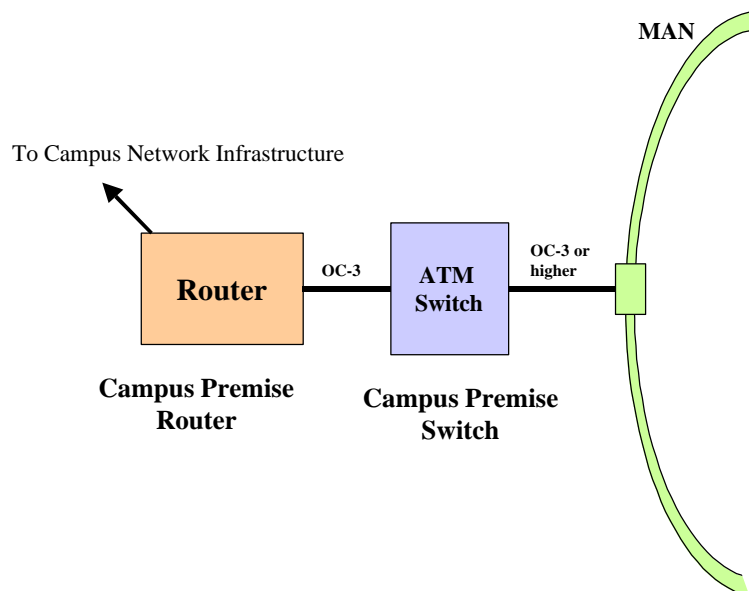


Figure 3-7. Campus Premise Router Connectivity to the MAN via Premise ATM Switch

The premise router needs connectivity to the other campus networks via the MAN. This should be done with a full Virtual Circuit (VC) mesh. The preferred choice is SVCs, which is accompanied by the challenge of figuring out how to perform the IP-to-NSAP mapping. Static mapping of the routers is one choice and probably the most stable, but another choice is RFC 1577 “classical IP.” The disadvantage of RFC 1577 is the instability of the Address Resolution Protocol (ARP) server as a single point of failure. Other choices are MPOA or IP Over NBMA (Non-Broadcast Multiple Access), which are both supported by a number of router vendors. Still another choice is to install a full PVC mesh as a single logical IP subnet (LIS) between the campus premise routers. This is a very stable configuration but is not very scalable – the number of resulting PVCs is on the order of N^2 (where N is the number of participating routers) and each requires significant manual configuration.

As a means of summarizing the above, the ITI IPT recommends that the preferred campus-to-MAN architecture be full mesh connectivity to the MAN via SVCs, static mapping of routers, and MPOA. Additional IP connectivity guidance is discussed in section 3.7.

3.4.3 Autonomous Networks

From an IP perspective, the “autonomous networks” each stand alone from other IP networks in the Naval enterprise. They have their own routing architecture and have their individual Interior Gateway Protocol (IGP) operated as a separate routing domain. The IP routers in an autonomous network are managed separately from the rest of the enterprise.

Autonomous networks can be connected to each other by using an appropriate protocol, e.g., Border Gateway Protocol (BGP) 4. The users connected to one autonomous network should be able to communicate with the users on another autonomous network. The separate autonomous networks are all logically part of the global Internet.

Each autonomous network will have its own security profile and may include security controls at its perimeter. Network managers should employ the appropriate security mechanisms, i.e. zone 3 or zone 4 protection in the defense-in-depth model.

This architecture recommends that any autonomous network that has a presence in a given region should perform peering using an Exterior BGP (E-BGP) mesh. This is similar to what is currently being done in the Washington, D.C., area in the National Capitol Region Metropolitan Area Network (NCR MAN). Traffic from one Naval autonomous network to another should use “hot potato” routing in which a packet is moved to the destination network as soon as possible because the destination network can best move the packet to its ultimate destination.

3.4.4 Fleet Intranet

The fleet intranet, which is a special case of autonomous networks, must be operated as a separate routing domain because of the highly dynamic nature of the fleet environment. In the fleet environment, ship communication links change often and when they do change, automation must replace human intervention to produce fast convergence of the routing service.

This architecture specifies routing the fleet IP traffic using a single Open Shortest Path First (OSPF) domain. The fleet support infrastructure ashore (NOC, piers, STEP sites) is considered the “backbone area” in the OSPF sense (also known as area 0). There is a separate OSPF area for each ship. Each ship announces a single aggregated prefix and those of any MEUs, air wings, or other deployed forces that are onboard. Large routing updates to the ships are collapsed into a default advertisement to save bandwidth. If additional hierarchy is required, a ship could be a separate routing domain (OSPF with multiple areas) and use BGP to peer with the “backbone” for redistribution.

Connectivity to the external world is accomplished through the NOCs/NCTAMS. They announce the aggregate fleet address/mask from outside the fleet firewalls. From outside the fleet intranet, it is impossible to determine ship locations. Traffic originating from the external world destined to the fleet must first go to the nearest fleet border (NOC firewall or some pier firewalls), and then traverses the fleet intranet until it arrives at its ultimate destination.

3.4.5 IP Addressing Plan

The goals of the addressing plan are to aggregate address space and to use address space efficiently.

The most obvious places for aggregation are at the campus and MAN levels. Addresses on campus should all be summarized with a single address/mask, and addresses for all MAN customers should aggregate to a single address/mask. There may be exceptions (legacy installations, unexpected growth, etc.), but these should be minimized. To perform this aggregation, the MAN operators should obtain a Classless Inter-Domain Routing (CIDR) block large enough to support the region and allocate sufficient address space to the campuses to

support the long-term needs of that campus. As a rule, a “class B” CIDR address block is adequate to support many campuses in a region and is inappropriately large for allocation to a single campus.

Unfortunately, large blocks of address space are difficult to acquire. An alternative is to use network address translation (NAT) technology at the campus or unit. With this alternative, non-unique address space is used internally but is translated into the globally unique space externally. This approach works quite well, and is sound from a security perspective. The NAT alternative is recommended whenever it is practical.

Many campuses in the DON already have allocated IP address space and do not want to change their IP addresses. Accordingly, these addresses in the near term will need to be globally routed.

A more prudent approach is for the Navy NIC to obtain a large block of address space and allocate it appropriately to each of the MANs. This would seem to be difficult initially, but once implemented would save significant management time.

To maximize efficiency of the address space usage, DON IT managers are encouraged to sub-allocate address space in appropriately sized blocks. For example, a point-to-point link should only be given a /30 (30 bits of mask, out of 32). A small IP subnet should only be given the amount it needs for the long term. Every entity should not automatically receive a “class C” chunk, (/24), unless it is justified to do so.

3.4.6 IP Routing Architecture

Figure 3-8 depicts the OSPF and BGP routing in the IP routing architecture. For the networks described, the preferred interior gateway protocol (IGP) is OSPF. Some autonomous networks may choose other protocols, but the MANs and campus networks will use OSPF. Each MAN will be configured as a “backbone area,” or area 0. Each campus on the MAN will be a separate OSPF area. The fleet intranet will also use OSPF.

BGP4 will be used between MANs and configured in a full E-BGP mesh. Because the WAN is implemented with a full VC mesh between the ITSC premise routers, there is no requirement for an additional IGP on the WAN. The BGP mesh is sufficient.

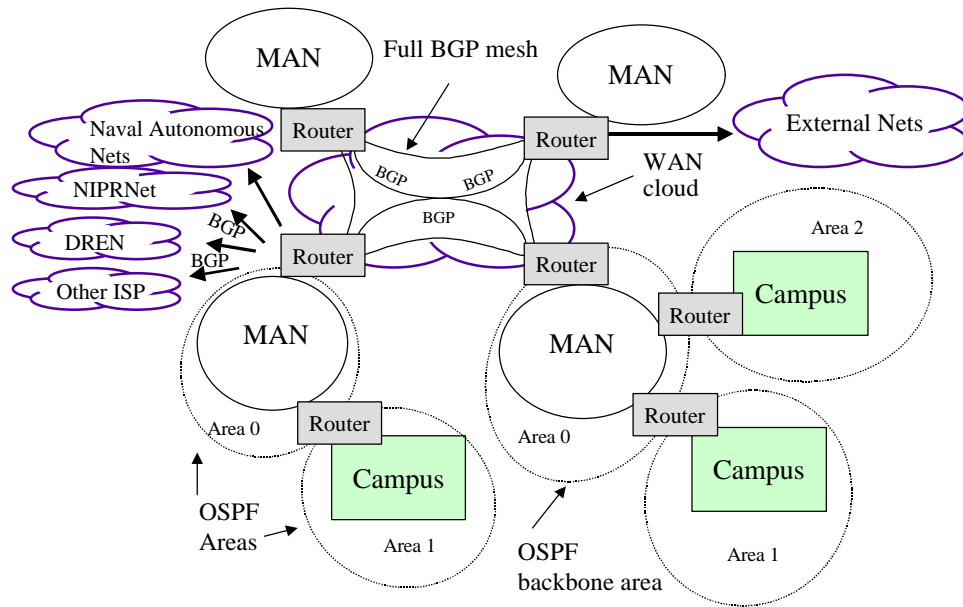


Figure 3-8. BGP and OSPF Routing in Enterprise Architecture

Between autonomous networks, BGP version 4 will be used. BGP4 will also be used to peer with the external world. Each autonomous network should only announce networks that are within its autonomous network. It should not announce networks that it learns from other BGP peers; it should not serve as a “transit” network. (There may be minor exceptions.)

Routing information between the IGP (OSPF) and the EGP (BGP4) should not be redistributed. Networks announced via BGP should be explicitly configured to ensure the highest network stability.

Naval autonomous IP networks should all peer at the ITSC to minimize routing distance.

3.5 Network Connectivity Security Overview

Defense in Depth is the preferred information protection approach for the DON. In Defense in Depth, information protection mechanisms are applied in multiple, complementary, and redundant locations to collectively form a system architecture. Defense in Depth is a layered approach analogous to the multiple zones around a carrier battle group. The layers provide maximum resistance to attack and minimize the likelihood that a single flaw can lead to a security compromise.

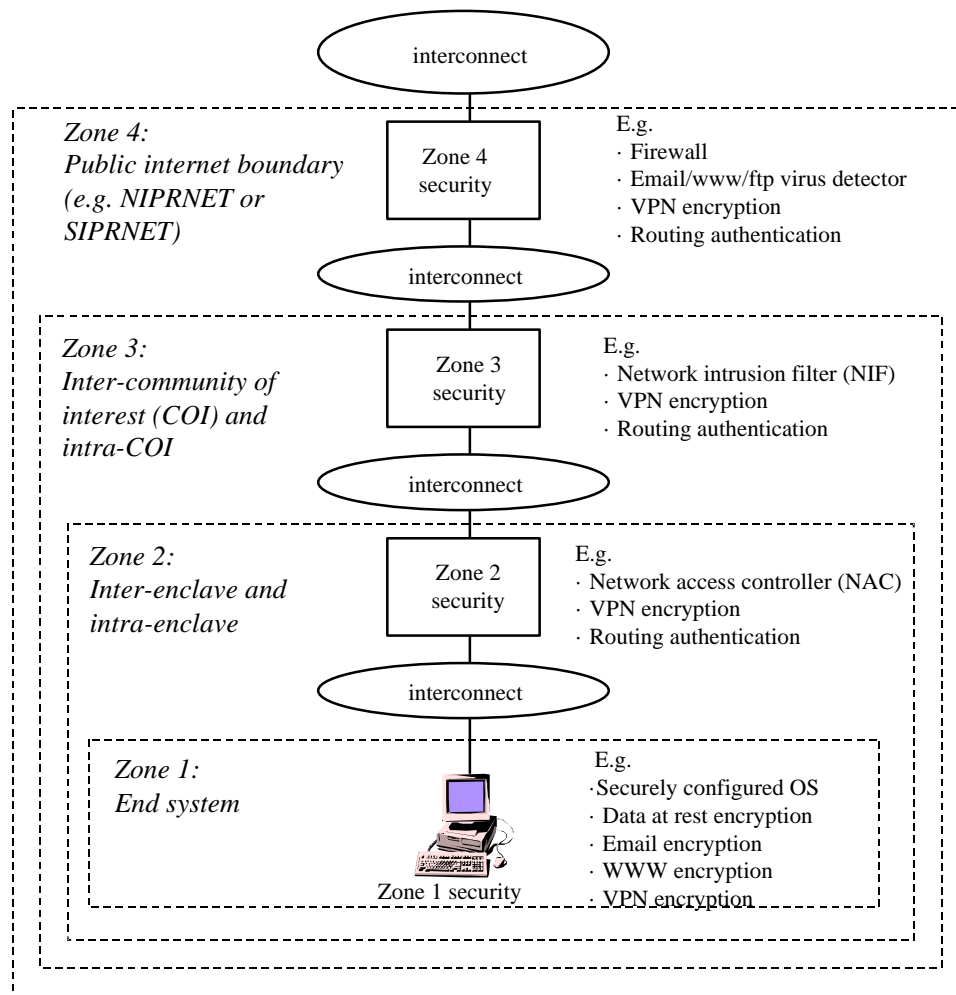


Figure 3-9. Generic Framework for Defense in Depth

A generic framework for defense in depth is illustrated in Figure 3-9. Four zones of defense are defined in this framework; these zones may be logical and not necessarily physically separate. Across the framework, numerous security mechanisms are used to protect information. In addition to the protection mechanisms, certain infrastructure components are required to build secure DON information systems. The most critical of these is a public key infrastructure (PKI) that provides identification and authentication mechanisms and encryption mechanisms for application across the various zones.

Appendix E provides a detailed account of the four security architecture zones and how the information protection components are applied to each zone and information dimension. Specific instances of how security is applied to area networks and network services is contained in each individual section of the document, for example, in this chapter – within ATM and IP, in Chapter 4 – within Directory and Domain Name System, and in the appendices – within each template.

3.6 ATM Connectivity - Detailed Architecture



As introduced in section 3.3.1, a successful implementation of ATM in the DON requires well-designed and coordinated planning across all Navy and Marine Corps organizations. Design factors to be resolved include address and routing strategies for interoperating in the joint DoD environment, service provider support of addressing schemes and ATM features/protocols such as Permanent Virtual Paths (PVPs) and Private Network-to-Network Interface (PNNI), and vendor support of ATM standards-based products. These and other implementation issues must be carefully evaluated in order to formulate an enterprise strategy that positions the DON to move toward the desired performance-oriented, cost savings potential of ATM.

The purpose of this section is two-fold. First, it outlines the ITI IPT's strategies and conclusions for choosing particular ATM architecture components. Second, it provides more detailed information on specific ATM protocols and addressing and routing plans.

3.6.1 ATM Planning and Implementation Constraints

The implementation is currently constrained by the following issues.

- Signaling (as opposed to tunneling) of virtual circuits over a public wide area carrier is difficult because service providers are not yet in full agreement on which standards to implement and when. For example, PNNI is not supported through the WAN but can be interfaced via Public UNI. Broadband InterCarrier Interconnect (B-ICI) is not implemented across service provider boundaries. Chosen solutions must be available to implement today and adaptable to future service-provider signaling offerings.
- ATM protocols are not only being introduced at an unprecedented pace but are still evolving. Chosen solutions must be fielded reliably today but yet take advantage of future, more robust protocols.
- Security solutions for ATM are still evolving. Basic NSA-approved Type 1 encryption exists today, e.g., the KG-75 Fastlane and the emerging KG-175 TACLANE, but traditional firewall-like devices are limited.

3.6.2 ATM Architecture Design Factors

Switched Circuits: End-to-end switched virtual circuits (SVCs) between any two DON ATM end systems is a requirement. Full signaling is sought from service providers but Permanent Virtual Paths (PVP) will suffice in the interim.

Ubiquitous ATM Service: Given the end-to-end SVC goal, it is desired that ATM service will be provided to every Naval base (campus), as appropriate.

Redundant MAN Links: No single switch or physical circuit will be a single point of failure for MAN access to the WAN.

Redundant CAN Links: No single switch or physical circuit should be a single point of failure for CAN access to the MAN (this requirement is driven by each particular site's mission).

3.6.3 ATM Addressing and Routing

3.6.3.1 Detailed DON ATM Addressing Plan

Implementation of the architecture described above requires a complete ATM addressing plan. The DON ATM Addressing Plan is based on the DoD ATM Addressing Plan as described in MIL-STD-188-176, Standardized Profile for Asynchronous Transfer Mode (ATM), dated 21 May 1996 and DISA implementation instructions contained in DoD ATM Addressing Plan, dated 17 April 1998.

The DON Network Information Center (NIC) is the single DON point of contact to the DoD NIC for both IP and Network Service Access Point (NSAP) registration. The NIC has established procedures to register ATM End System Addresses (AESAs).

DISA assigns the routing domain fields of the ATM address and the DON NIC assigns area field values to DON regions. Because DISA assigns the routing domain field (bytes 5 through 8) geographically, each Navy MAN will have a different address prefix.

For fixed sites, the NSAP addressing will be geographically based using the DISA NSAP addressing plan (as modified for DON MANs).

For the fleet autonomous networks and components, home-port geographically-based addresses will be used. A given autonomous network/component will aggregate to a single prefix/mask that will follow the unit; re-addressing of network devices is not acceptable. The network itself will adapt to the change in component location. Further, nesting of components will be supported, e.g., an MEU could deploy from its home base to a ship that in turn deploys to another area — requiring no change in addressing of either.

The System ID portion of the address range is unique within the MAN.

All addresses on a given campus will be summarized by a single prefix/mask and advertised to the MAN. Also, all addresses within a geographic area covered by a MAN and the campuses it

serves will be summarized into a single prefix/mask and advertised to the WAN. This hierarchy makes the routing problem more tractable.

An exception to the above is a campus that is dual-homed to a non-Naval ATM network. In this case, the campus may have to follow the addressing conventions of another and request that all providers accept those addresses for networks on which signaling support is required. This may be difficult when one service provider is unwilling to carry another service provider's NSAP address. Should this case arise, the dual-homed site must choose an NSAP address that is acceptable to all service providers. (DISA, for example, has agreed to carry non-DISA addresses on the DISN on a non-operational case-by-case basis.)

The scheme for fleet address advertisement is still being designed. Some of the constraints are to not require re-addressing of network components, minimize and hopefully eliminate the need to manually propagate static routes, not disclose the state or location of deployed forces, and allow for efficient routing between warfighting elements.

The DISA ATM addressing allocations provide a 24 bit (3 byte) field called "Area" for subscribers to use in assigning NSAP prefixes to switches. These would be applied at the MAN level. The MAN will then allocate address space down to the individual sites (campus networks). Figure 3-10 gives an example of the allocation hierarchy that can be used for a typical MAN.

| | | | |
|----------|--------------------|--------------|------------|
| Site (6) | classification (2) | Building (8) | Switch (8) |
|----------|--------------------|--------------|------------|

Figure 3-10. Campus ATM Address Allocation (without Pier)

The allocated 24 bits are used to establish an addressing hierarchy in a region. Six bits allows 64 distinct sites to be connected to the MAN. Two bits of classification allows for four different classification levels at any given site. Eight bits allow for 256 buildings at a site. And finally, eight bits allow for 256 switches in a given building.

There is substantial latitude for variation in this allocation hierarchy. A given site can adjust its share of the suggested address space to meet local requirements. For example, if the MAN takes six bits for site designator, the site can allocate the remaining 18 bits, depending on local requirements. A site with a few large buildings may have a very different allocation scheme than a site with hundreds of small buildings spread over a large area.

Another variation is at the MAN level. If there is a small number of large sites in the region, along with a large number of small sites, then a dual allocation scheme could be used. One example is eight bits of site prefix for small sites and three bits for large sites. Thus, there would be $8 + 256 = 264$ sites total where a large site would have 21 bits to assign instead of 18.

Sites with piers for ship home-porting have additional requirements and are represented in Figure 3-11. Because the ship ATM prefixes are globally routed, each ship requires its own globally-unique prefix. From a campus network perspective, the ship can be considered to be a building. However, at the shipboard level there are multiple classification levels, so the bit allocation must be adjusted to meet that requirement.

| | | | | |
|----------|--------------------|----------|--------------------|------------|
| Site (6) | classification (2) | Ship (8) | classification (2) | Switch (6) |
|----------|--------------------|----------|--------------------|------------|

Figure 3-11. Campus ATM Address Allocation (with pier)

This allocation supports a site with up to 256 ships home-ported. This is more than enough, because the largest homeport is currently assigned 81 ships. This scheme also allows for four classification levels and up to 64 switches per classification level. Variations are supported as previously described.

3.6.3.2 Detailed DON ATM Routing Architecture

The MAN prefix/mask aggregate must be “advertised” within the WAN. Depending on the WAN provider, this might be done any number of ways. A fully-signaled, automatically-advertised solution such as full PNNI is strongly desired. Given the immobile nature of MANs, however, static addressing through Public UNI or a PVP mesh is acceptable in the short term.

For the fleet, the ashore components of the fleet autonomous ATM network will use PNNI routing and will be a peer group at the top of an independent routing hierarchy. These components will be part of a two-level PNNI hierarchy with ships at the bottom and the ashore infrastructure at the top. Each ship will be a logical group node, but each group can also be a logical node, which allows it all to collapse into a single level. Additional levels can be added at the appropriate time (battle group, theater). This architecture will be very simple and will only create additional hierarchy when it makes sense.

The simpler implementation is shown in the flat hierarchical routing depicted in Figure 3-12.

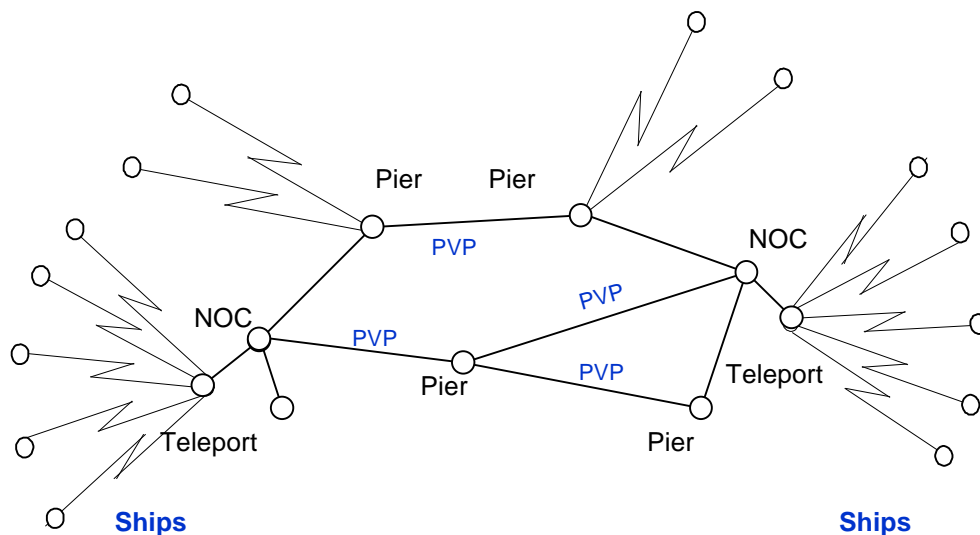


Figure 3-12. Flat PNNI Routing

A far more complex and long-term implementation is shown in the deep hierarchical routing depicted in Figure 3-13.

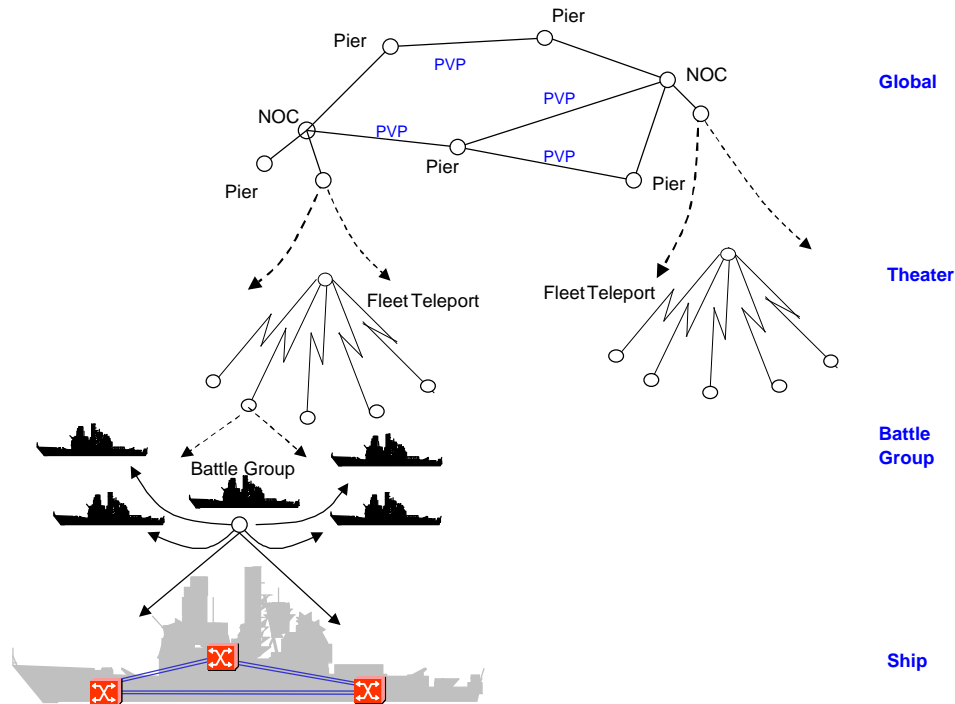


Figure 3-13. Deep PNNI Routing Hierarchy

While a deep routing hierarchy like this is possible, it is unclear whether it will be necessary or even beneficial. The advantage comes only if routing topologies can truly be summarized at the levels shown above. Unfortunately, the fleet is not static or predictable. Since the additional complexity offers little or no advantage, the simple flat hierarchy makes the most sense for the foreseeable future.

Whether a ship is connected via a pier or via an ATM-based satellite link, it will announce its ship prefix/mask via PNNI to the fleet autonomous network (ATM PVP overlay). The fleet network will then have full routing information for all ships, so signaling will take a direct path. When signaling originates outside the fleet network, the call will be routed to the region where the ship is home-ported. At that point, it will go directly to the ship if it is pier-side or will be routed to the fleet network if it is deployed.

3.6.4 ATM Protocols

A brief description of selected ATM protocols is contained in this section. More detailed information is available at the ATM Forum web site (<http://www.atmforum.com/>).

3.6.4.1 Private Network-to-Network Interface (PNNI)

This section on PNNI is provided for consideration by network planners. PNNI is critical to realizing the intended performance, interoperability, and ease-of-use of ATM.

PNNI is a dynamic routing protocol defined by the ATM Forum and specified for use between ATM switches and between groups of ATM switches. PNNI includes two categories of protocols.

- **Routing.** The ATM PNNI protocol allows each switch in the network to share topology information. This information is used to determine the best paths for an end-to-end route through the ATM network based on parameters such as Quality of Service (QoS) and advertised bandwidth use. A key feature of the PNNI mechanism is its ability to automatically configure itself in networks in which the ATM address structure reflects the hierarchical routing topology.
- **Signaling.** PNNI also defines a network-component protocol for signaling or message flows used to establish point-to-point connections across the ATM network. This protocol is based on the ATM Forum UNI signaling and uses the routing services of PNNI Routing to provide optimized end-to-end path selection for user calls.

A hierarchical relationship among ATM switch addresses is essential to successful operation of the DON Technology Infrastructure. Besides supporting a logical ATM addressing structure based on the network topology, the hierarchy makes ATM call routing simpler and more efficient. The hierarchy also enables scalability to a large world-wide ATM network.

The hierarchical addressing structure differs from the flat approach traditionally used in IP routers. In a flat address space, every router must maintain information about the topology of the entire network (or at least a “default” path to a more knowledgeable router). This results in excessively large routing tables and inefficiency. Networks based upon the first three bits of an IP address, together with a subnet mask, was presented as the original solution to this problem. Later, the notion of autonomous system numbers and the use of different routing protocols within IGP and between EGP networks was employed to contain the growing size of routing tables. Finally, classless interdomain routing was introduced primarily to conserve IP addressees and to further contain the growth of routing tables. In a hierarchical address structure such as PNNI, routing information is aggregated when detailed knowledge of the topology is not needed and a single routing protocol is used throughout the network.

The DON ATM hierarchy is created through the assignment of Network Services Access Point (NSAP) addresses. The network uses longer unique prefixes of the NSAP address for more detailed routing information (lower in the PNNI hierarchy). With the bits available in the NSAP address, there are levels in the hierarchy to enable efficient routing.

Each DON enterprise switch must have a unique address prefix/mask which is different from the address prefix/mask assigned to any other switch or device in the same network. All subordinate switches and ATM devices must have addresses based on the primary switch in the hierarchy.

Hierarchical addressing is the point of the PNNI Phase 1 (PNNI-1). The PNNI-1 protocol supports the concept of peer groups. Each peer group consists of multiple ATM switches which operate in the same hierarchical level. They communicate through a peer group leader who represents the peer group at the next higher layer of the hierarchy.

It may be necessary to establish PVPs between Navy and Marine Corps regions to facilitate top-layer PNNI routing within the DON. These PVPs might not be necessary if the architecture consisted of a collection of MANs using DISA (or an alternative WAN ATM service provider) as the communication infrastructure for interconnectivity between MANs. The goal of the DON enterprise backbone services is to provide a single PNNI peer group hierarchy.

3.6.4.2 Multi-Protocol over ATM (MPOA)

MPOA is the protocol for interoperability of both ATM and non-ATM devices in an ATM environment. While not emphasized in the early stages of this architecture strategy, MPOA will become increasingly important to improve DON network performance.

MPOA is the first standards-based protocol that allows routed networks to take advantage of the benefits of the ATM network (i.e., lower latency, performance, and scalability). It expands on schemes such as LAN Emulation, Classical IP/ATM (RFC 1577), and Next Hop Resolution Protocol (NHRP) to create a standardized notion of a virtual routing functionality integrated within a high-speed, dynamically-switched ATM network. MPOA, by allowing traffic to be forwarded to its destination over an ATM virtual circuit, reduces the cumulative latency in a multi-protocol routed network by reducing the number of intermediate points where packet processing must be performed. MPOA also allows non-ATM network layer protocols to take advantage of the QoS features of ATM.

The MPOA “switched-routing” methodology for IP consists of the following components:

- **ATM Network Cards/Drivers** allow ATM directly-connected hosts to send and receive traditional IP datagrams to and from an MPOA-capable network and interoperate with non-ATM hosts which are indirectly connected.
- **IP Switches** integrate routing (layer three) intelligence into an inherently switched (layer two) transport infrastructure (hence switched routing). The IP switches perform the packet-forwarding function on the non-ATM edge devices or hosts in an MPOA network.
- **Route Servers** maintain MAC address routing tables and act as distributed directories to translate the destination MAC addresses to the ATM address of the destination switch. Once the destination ATM address is known, the source device can establish a VCC directly to the destination device. This “cut-through routing” has the performance impact of eliminating all but a single IP “hop” from any source to any destination within the DON enterprise network. The resultant IP “diameter” of the network is 1, which is the smallest diameter possible. Participation in “cut through routing” is controllable, i.e., designers can choose which ATM devices are or are not allowed to accept “direct” VCCs through policy or flow patterns. The MPOA specification defines a virtual routing framework that separates the route calculation function from the actual layer 3 forwarding function. The MPOA capable edge device or host provides layer 3 forwarding of packets.

MPOA enables improved performance and interoperability of both ATM and non-ATM networks in wide area ATM network domains, not only in the MAN template environment but also in the WAN environment. It has even greater implications for the campus LAN environment.

MPOA is a logical evolution of ATM LANE and is upwardly compatible, i.e., a LANE host/edge device can participate in an MPOA-capable network.

3.6.4.2.1 Design Factors

MPOA is required for implementation in ATM networks based on this architecture guidance and the Information Technology Standards Guidance.

The WAN must support MPOA for TCP/IP.

3.6.4.3 LAN Emulation (LANE)

ATM LANs, non-ATM attached LANs, and accompanying end stations need to communicate over ATM networks. Prior to the deployment of MPOA, LANE was the acceptable technique to enable communication in this mixed environment. This is possible because the ATM network "emulates" the characteristics of broadcast LANs (e.g., Ethernet, FDDI, and Token Ring) and performs Media Access Control (MAC)-to-ATM address resolution. Implementation of LANE consists of a LAN Emulation Server (LES), a Broadcast and Unknown Server (BUS), and a LANE client (LEC). Each component resides in one or more ATM end systems or edge devices. Although it is possible to implement these functions in an ATM switch, this is usually avoided (as explained below). The LES and BUS work with LECs, typically Ethernet hubs with an ATM uplink, to provide layer 3 bridging functionality across the ATM network along with directly attached ATM hosts.

While it is possible to provide support for legacy networks such as Ethernet or Token Ring via LANE services in the core, most LAN emulation will be implemented in the local edge switches and ATM-connected hosts.

An edge switch is on the boundary of an ATM network. Typically it is an ATM switch which supports legacy networks via LANE services or an ATM end system workstation/host. For smaller networks, the edge switch can connect directly to the core.

The decision as to where to run LANE depends on the particular ATM devices that handle the LANE processes. For autonomous networks and/or communities of interest, the LES/BUS services should be implemented on ATM end systems (or on the edge switch) close to the units that are connecting to the various emulated LANs (e-LANs). (If desired, redundant LES/BUS servers should be employed.) In addition to required redundant hardware, there need only be one LANE Configuration Server (LECS) per ATM network.

Every e-LAN functions independently by using its own LES/BUS. Interconnecting multiple e-LANs requires a router typically called a one-armed router, named so because there is only one physical interface (usually 155-Mbps OC3 fiber) from the ATM switch to the router. Instead of routing between multiple physical interfaces, the one-arm router simply routes the layer 3 protocols onto multiple e-LANs connected through the same interface. The router can be a card that is inserted into an ATM switch or it can be a stand-alone router.

The natural upgrade path from LANE is MPOA.

3.6.4.4 Voice and Telephony over ATM (VTOA)

VTOA allows a traditional voice circuit to be dynamically signaled and mapped into an ATM virtual circuit (as opposed to circuit emulation in which voice circuit trunks are provisioned — not signaled — through ATM networks). VTOA provides improved bandwidth effectiveness because the voice calls are created and terminated on demand over the same network that is transporting data. VTOA also assures voice quality by employing voice adaptation techniques and Quality of Service support. VTOA reduces cost by consolidating networks for voice, data, video, and imagery via an ATM network.

WAN voice communications have traditionally used analog and digital leased lines between major locations supported by PBX switches or by public Centrex services using tandem PBXs to minimize facility costs. The design of PBX networks has changed little over time and produces poor efficiency (e.g., bandwidth use, compression degradation, blocked calls) and effectiveness. While circuit emulation can be used in the near term by provisioning traditional voice trunks through a network, end-to-end signaling is preferred. The DON enterprise architectural strategy is to provide MAN services that fully support the consolidation and implementation of voice. The general use of distributed sets of Remote Switching Modules (RSMs) that are homed to central office switches provides a Service Delivery Node. This supports use of existing infrastructure and centralization of egress points to minimize circuit costs, centralize trouble reporting operations, and expand the network to encompass additional users and locations.

The DON TI architecture for voice includes two options for the interworking between Defense Switch Network (DSN) and the ATM WAN:

- Based on Circuit Emulation Services (CES) in the ATM WAN that is transparent to the DSN.
- Based on PVC or SVC that map the DSN signaling and user traffic to and from ATM formats to provide cell-based direct interworking between the DSN and the ATM network. In a PVC architecture, the VTOA IWF and ATM infrastructure are transparent to the DSN switching systems. Calls are routed to the destination DSN switching systems over an ATM virtual circuit trunk carrying the narrowband traffic and associated signaling. In an SVC architecture, the VTOA IWF must process a subset of the DSN signaling messages to establish SVCs to carry the DSN calls.

3.6.5 ATM Overlay Security

The DON ATM overlay shall provide a private or virtual private network for the DON. The DON ATM overlay must satisfy the following security requirements:

- Confidentiality
- Integrity
- Reliability
- Robustness

The DON ATM overlay must provide data transport service at the secret (S) and sensitive but unclassified (SBU) system high levels.

It is assumed that CANs (switches, routers, fiber, etc.) are under positive DON control and are operated at the SBU level. It is also assumed that MANs and WANs may NOT be under positive DON control and are operated at the unclassified (U) level.

In order for the ATM overlay to satisfy confidentiality and integrity requirements, ATM cell encryption shall be used. For secret data transport services, information shall be encrypted using KG-75 Fastlanes before entering a DON CAN or Zone. For SBU data transport services, information shall be encrypted using KG-75 Fastlanes before entering the portion of the ATM overlay that is NOT under positive DON control. For example, if a MAN is constructed by leasing ATM service from a commercial ATM provider, this encryption would be applied at the CAN/MAN boundary.

In order for the ATM overlay to satisfy the requirement for reliability and robustness of data transport services, the network must provide redundancy, guarantee the minimum bandwidth required to support high priority data transport, and provide mechanisms for managing network bandwidth allocation (particularly when contention occurs).

Redundancy shall be such that no single failure of a network component or interconnection leads to the isolation of any CAN from the overall DON enterprise network. The careful satisfaction of the redundancy requirement will require detailed analysis of external vendor network configurations when commercial ATM services are used to implement MANs and WANs.

Guarantees of minimum bandwidth can be provided in a number of ways. These include physically dedicating bandwidth using either dark fiber (leased or owned) or SONET channels, and PVPs with appropriate committed information rates. In addition, if DISN ATM services are subscribed, it may be possible to use SVCs when MOAs guaranteeing minimum bandwidth can be established.

Management of bandwidth allocation within the DON ATM overlay must be provided. This management must provide authorized administrators with the capability to identify the priority of data transport requests and allocate network bandwidth to the highest priority transport requests when contention occurs. In addition, the ATM overlay must feature mechanisms to order and add additional bandwidth as required.

Finally, the ATM overlay must provide mechanisms to ensure that DON-controlled components of the overlay can only be managed by authorized administrators, are resistant to penetration attempts, and are resistant to ATM signaling-based denial of service (DoS) attacks. The use of KG-75 Fastlanes significantly reduces or eliminates the potential for unauthorized administration and successful penetration originating from outside the DON-controlled portion of the overlay. In order to reduce the potential of successful penetrations originating from inside the DON-controlled portion of the overlay, network components that are remotely managed must feature a non-spoofable authentication mechanism. Use of KG-75 Fastlanes may reduce the potential for successful ATM signaling DoS attacks originating from outside the DON portion of the overlay. However, due to the current lack of authentication in ATM signaling protocols, it is difficult to protect against attacks originating from inside the DON-controlled portion of the overlay.

A comprehensive treatment of ATM security objectives and threats is well documented in the ATM Security Framework 1.0 (ATM Forum AF-SEC-0096.000, Feb. 1998) and parallel those of IP networks. Efforts are underway at DoD to resolve these ATM security issues.

3.6.5.1 Design Factors

ATM implementations must reflect the most up-to-date ATM security guidance.

Encryption for ATM is provided through the use of Fastlanes or TACLANEs. These devices encrypt the data payload of ATM cells but not the header. When traffic analysis is a concern, link encryption devices may also be used.

Remote management of ATM devices is a concern. SNMP should not be used for anything other than collecting status information. Remote login to these devices requires non-spoofable authentication such as token-based access control. Whenever possible, SNMP access and traffic should be limited to in-band, non-routable subnets.

Regarding denial of service, signaling in ATM is a special concern. ATM signaling does not have authentication. Guidance on safe ways to use signaling in ATM will be provided at a later date.

3.7 IP Connectivity Detailed Architecture

As introduced in section 3.3.2, the IP addressing architecture is focused on the myriad organizational and operational characteristics of the Naval enterprise. For example, commands vary in size, level of complexity, and networking knowledge and are distributed worldwide. Ships deploy and transit between theaters while maintaining network connectivity so that a ship's point of presence on the network changes constantly as its location changes. The latter condition is exacerbated by relatively low-bandwidth, high latency, and high error rates, which are all common with satellite and line-of-sight radio frequency (RF) voice and data links.

The network architecture comprises the connectivity upon which traffic is transported throughout the Navy and Marine Corps. The connectivity spans resources and activities across the physical, network, and application layers and uses a mixture of ATM and IP technologies. Combined, these two technologies allow for the creation of a high performance, flexible, and proven infrastructure.

This network connectivity uses ATM standards to create flexible, high-bandwidth, IP-based "autonomous networks" on top of the ATM "virtual circuits". This design supports flexibility to provision bandwidth where needed without being hampered by congested IP paths or too many router hops (i.e. network diameter). This services profile extends to all regional ITSCs and to each campus. In support of the fleet intranet, it extends to the piers, NCTAMS, and STEP sites.

At higher networking layers, voice, video, and data network services are provisioned. For data, this infrastructure is established through multiple IP subnets. These "intranets" consist of logically-separated autonomous networks that are part of the larger DON enterprise network and which, for security reasons, require some level of logical segmentation. This architecture provides for the creation of these virtual networks for both large commands and concentrations of geographically close smaller commands that obtain services from a MAN. The separate IP subnetworks are constructed on top of the ATM services provided in the network layer below

them. These IP subnetworks are aligned with separate security enclaves or autonomous networks that manage their own routing domains.

3.7.1 IP Connectivity Design Factors

The DON enterprise network does not carry IP traffic between other networks if neither the source nor the destination of the IP traffic is a DON enterprise subnetwork host. Exceptions will be made for legacy DON IP networks which will be treated at ingress and egress as external (i.e., Internet) addresses and will be subject to Internet access control policies.

Unrestricted IP traffic between regions (i.e., the associated IP networks) is supported on the backbone using methods described in the previous sections.

Access to DON enterprise network resources from legacy DON IP networks is controlled by Zone 4 security protection mechanisms.

In order to minimize the size of internal routing tables, regions obtain a CIDR block sufficiently large to provide IP service to all tenant organizations within the region. This is based upon the number of campuses and the population of each campus.

3.7.2 Placement of Routers in IP Connectivity

The connectivity of campus routers to a MAN via a full Virtual Circuit (VC) mesh was described in section 3.4.2.

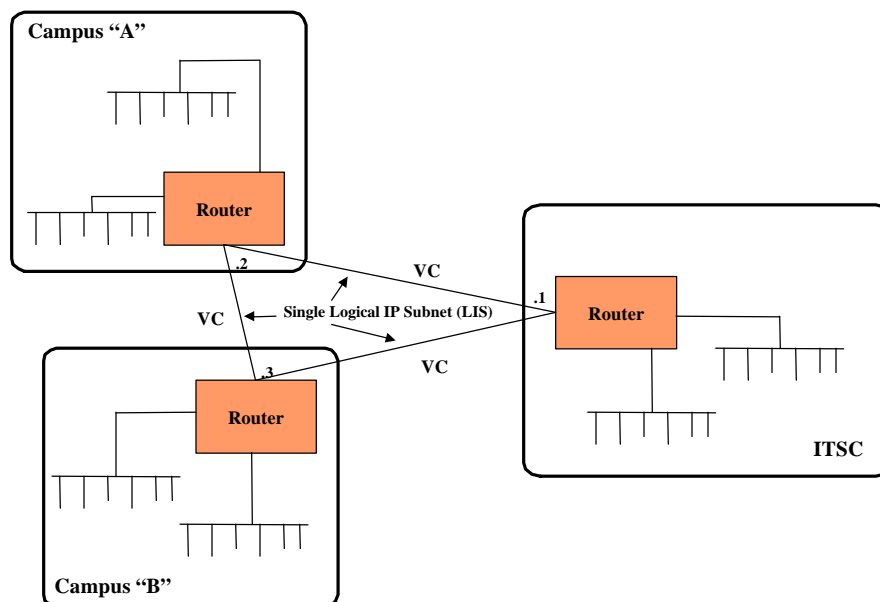


Figure 3-14. Multiple Campus Premise Routers, Showing VC Mesh Between Them, with Single LIS

One of the participating routers in a connectivity mesh will be the ITSC premise router. As shown in Figure 3-14, this router must include connectivity to the other Naval MANs and to the external world (NIPRnet, Internet, etc.) via the ITSC firewall. For reliability purposes, there should be two parallel routers; the figure is simplified for discussion purposes. To reach the other Naval MANs, there is a VC mesh over the WAN using the same connectivity principles as the MAN. This links all ITSC premise routers. The same recommendations and tradeoffs that are in the MAN case apply here.

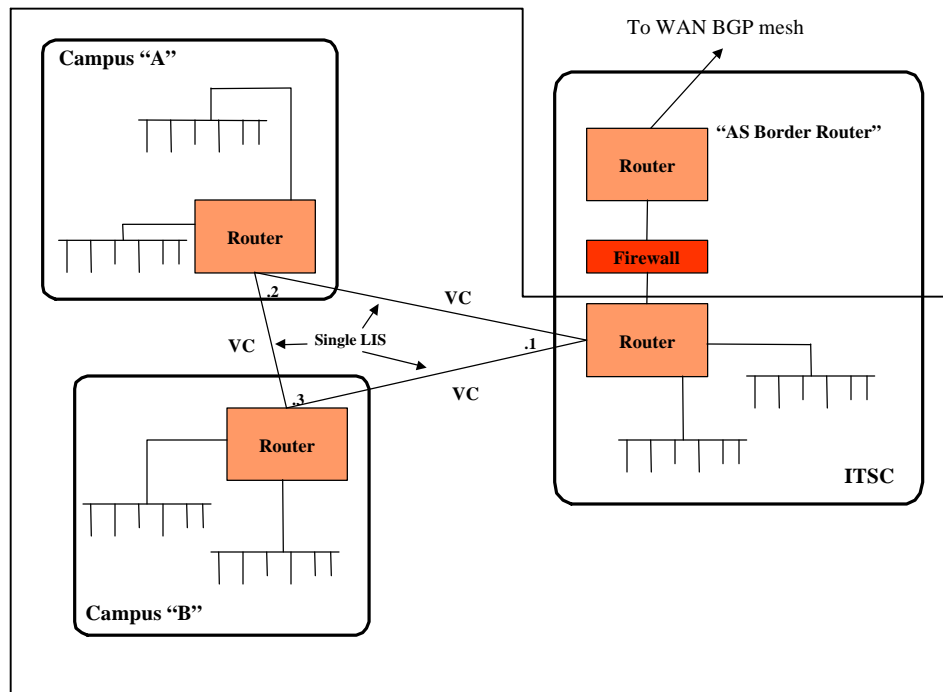


Figure 3-15. Multiple ITSC Premise Routers, Showing Full VC Mesh Between Them, Plus the ITSC Firewall, with Connection to the Outside World

There are cases in which a campus needs connectivity to the ITSC outside the firewall. Figure 3-15 shows that possible reasons for this may include an interim connection while back-door connections are eliminated. It may also be that the campus has a research network that needs to stay outside the firewall because of connections to external research nets. It may also be that the ITSC needs to peer with the external connections available at other campuses.

To accommodate these cases, it is clear that a separate VC mesh is required between routers that must peer outside the ITSC firewall. The architecture will be identical to the internal mesh described above.

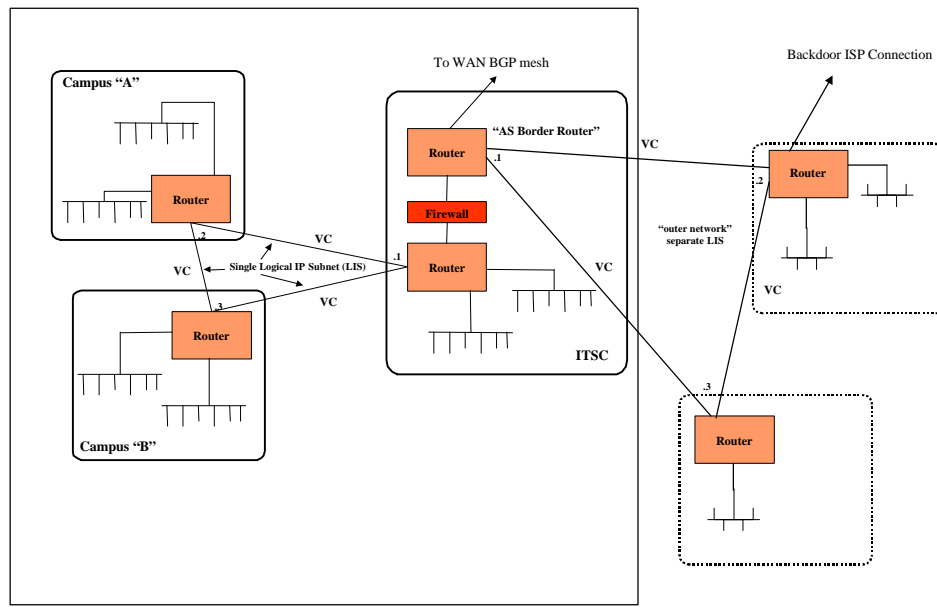


Figure 3-16. Picture of Multiple Campus Premise Routers, with VC Mesh, Connected Outside the Firewall

In order to minimize the number of routing hops between campuses and regions, the next step may be to implement MPOA to provide cut-through services for IP over ATM. It will not be necessary to do this right away, but this performance optimization can be designed and implemented at the appropriate time. Traditional router functionality, i.e., filtering, can be done on an MPOA-capable router. In this case, a cut-through virtual circuit will only be established if the source/destination IP addresses are “approved” to communicate.

The important point above is that all campus premise routers have direct access (zero hops) to the other campus premise routes on the MAN. Also, the ITSC premise routers all have direct access to each other via the WAN.

LANE technology should not be used across the MAN or WAN but should be restricted to the campus level.

3.7.3 IP Addressing

Regional commanders obtain IP network address space on behalf of ships and shore commands within their region from the Naval Computer and Telecommunications Station (NCTS) Pensacola, Florida, at (COMM) 850-452-3501, (DSN) 922-3501, or on-line at the Navy IP Network Number Registration page at <http://www.netreg.navy.mil>.

3.7.4 IP Routing

Establishing an appropriate routing architecture is critical to the success of the DON enterprise WAN. The connectivity, routing, and addressing schemas between ship, shore, MAN, and campus infrastructures are appropriately designed so that the enterprise network is scaleable and the number of hops between endpoints is minimized.

The routing architecture to address the fleet environment consists of a single routing domain for the entire fleet. Open Shortest Path First (OSPF) is used as the interior gateway protocol (IGP) for routing. The combination of dedicated and virtual circuits employed below layer 3 provides the networking infrastructure to support this design. The ITSCs and piers are included in this routing scheme. Ashore locations provide connectivity to the rest of the DON backbone network and the Internet and provide firewalls for Zones 2 and 4. Additionally, using OSPF allows fast convergence in response to topology changes that occur more often in the fleet environment.

Each of the MANs will need to obtain an Autonomous System (AS) number. This is required to support BGP peering with networks outside the region and for campus networks that already have their own ASN. Any given network should only be sourced from a single ASN.

The design factors for routing include determining required support of the following: LANE, Classical IP and ARP, and MPOA.

3.7.5 IP Overlay Security

The DON IP overlay shall provide a private or virtual private IP network for the DON and provide controlled interconnections to external IP networks. The DON IP overlay must satisfy the following security requirements:

- Confidentiality
- Integrity
- Reliability
- Robustness

The DON IP overlay must be provided at the secret and SBU system high levels.

The secret DON IP overlay will provide controlled interconnections to the SIPRNET and the SBU DON IP overlay will provide controlled interconnections to the NIPRNET. The interconnections must provide a high degree of isolation and penetration resistance from the NIPRNET and SIPRNET while allowing the flexible interchange of data between subscribers to the DON IP overlay. In addition, the DON IP overlays must continue to function in the event of SIPRNET or NIPRNET failure.

The DON IP overlay requirements for confidentiality and integrity shall be satisfied by using the DON ATM overlay at the appropriate system high classification level (secret or SBU). Also, the DON IP overlay requirement for reliability and robustness may be partially satisfied by virtue of the assured data transport provided by the DON ATM overlay.

The DON IP overlay requirements for privacy and controlled, secure interconnection to SIPRNET and NIPRNET (including continued operations of the DON IP overlay in the event of a NIPRNET or SIPRNET failure) shall be satisfied by employing a defense in depth approach to security. Key elements of this approach include network firewalls, secure domain name service (DNS), public key infrastructure (PKI), network intrusion detection, content security checking, secure configuration of workstations and servers, remote management security, and routing protocol authentication. Optional elements of this approach include data at rest encryption, virtual

private network (VPN) encryption, host-based intrusion detection, network intrusion filters (NIFs), and network access controllers (NACs). Detailed information on all of the elements can be found in appendix E.

From an architectural standpoint, network firewalls must be given particular attention in the DON IP overlay. Network firewalls shall be located between the DON IP overlay (Secret and SBU) and external IP networks (SIPRNET and NIPRNET). These firewalls, when combined with the other elements of the defense in depth approach, provide the controlled, secure interconnections to SIPRNET and NIPRNET. In addition, the combination of these firewalls and the use of the inherently robust/reliable DON ATM overlay provide a basis for satisfying the DON IP overlay requirement for robustness and reliability.

3.7.6 Considerations for Connecting Contractors

The customer frequently has a requirement to connect contractor and other non-Naval networks directly onto the Naval infrastructure. The reasons generally fall into two basic categories:

- Unsatisfactory performance of an Internet connection between a contractor site and a Naval site. This occurs because the contractor site receives connectivity via its own corporate intranet which has a gateway to the Internet. Traffic from the contractor site transits through their corporate infrastructure, through the network peering points, through the NIPRnet, and finally to the Naval site. The traffic may incur many router hops and congestion, even where the end points are at the same location. A direct connection between the two users would greatly improve performance.
- The contractor site is blocked access to the Naval site. The contractor appears as an Internet connection that cannot be trusted and is blocked for some access levels. To communicate, the contractor needs to bypass a security perimeter, such as a firewall or router filter, and establish direct connectivity “inside” the Naval infrastructure. To penetrate the security perimeter and gain access to systems on the protected side of the network, it must be done so that it does not compromise the security architecture.

If the main issue is performance and the connection to the Naval infrastructure is established outside the security perimeter (firewall or otherwise), then no additional security is required. If, however, the requirement is to bypass the security infrastructure, then it must be done in a manner that does not weaken or compromise the overall security architecture.

When non-Naval networks connect to the Naval Intranet inside the firewall, the following issues must be addressed:

- Does the contractor site connect to other networks? If so, these will constitute back-door connections and the connection should not be allowed.
- Can the contractor network be segmented with one segment being a separate network having no back-door connections? If so, this provides adequate protection and enables satisfactory connection to the Naval infrastructure. This segmentation may not be a desirable option for the contractor site.
- Does the contractor site have adequate controls of physical access? For example, are there locked doors? Who is allowed access to the protected network? What are the policies and

procedures for establishing a connection to the protected network? Who can use the workstations that are connected to the physical network?

Extending the Naval Intranet to a contractor site opens up many new opportunities for security compromise because the contractor site is outside the Naval physical security perimeter. An acceptable solution must satisfactorily address the requirements and issues described.

The Naval Intranet will provide external access using Virtual Private Network (VPN) technology. This allows access to the Naval Intranet from anywhere on the global Internet by establishing a secure tunnel through the Internet to the Naval Intranet. Access can be controlled by use of an identity certificate provided by the Naval/DoD Public Key Infrastructure (PKI) solution described in the next chapter. Using PKI, access to the Naval Intranet is controlled on an individual basis instead of on the network or device level. Access to the Naval Intranet is granted only to those individuals who have a recognized need for network access. Authorized users must still be authenticated to gain access to systems and applications on the network, just like any other Naval user.

This VPN solution provides a general solution for contractor access, as well as for anyone (including Naval personnel that are travelers and telecommuters) who requires access from outside of the Naval Intranet. In this manner, direct connections can be established from outside the security perimeter while preserving the integrity of the security architecture.

3.8 DON ITI Architecture Plan of Action

This section reflects the priorities and steps that should be addressed in developing an ITI architecture and for planning and implementing selected components such as MANs and CANs.

3.8.1 Steps for Developing Detailed Enterprise ITI Architecture

The following outline reflects the steps that the ITI IPT will use to fully develop the ATM detailed architecture.

1. Identify the MANs and major sites in the Navy and Marine Corps that need ATM. The initial recommendation of the IPT is that every MAN and most campuses should have ATM-provisioned service.
2. Determine the list of components to be connected to ATM. A list of questions is relevant to making this determination:
 - Which of the components require ATM?
 - Which of the components can operate effectively with IP?
 - How much bandwidth is required?
 - What is the existing external IP and/or ATM connectivity?

These questions should be posed for the following:

- Major components
 - ♦ MANs
 - ♦ ITSCs
 - ♦ Piers
 - ♦ Teleports
 - Intermediate components
 - ♦ Large campuses not near a MAN
 - Small components
 - ♦ All the small bases
3. Develop a high level global wiring diagram that shows the world-wide connectivity required for the Naval enterprise network, including the above components.
 4. The conclusion of the IPT is that a multi-level PNNI hierarchy (WAN/MAN/Campus) is not a pragmatic solution for the near-term. This is in consideration of the need to provide a Naval enterprise solution for global routing and addressing across ashore and afloat platforms in view of the immaturity of standards and conventions for dynamic routing across public carriers. The near-term target architecture configuration for which the IPT will provide a clear definition during the coming months is as follows:
 - WANs: PVP service, end-to-end signaling
 - MANs: If SVC service available, fully supported signaling environment and PNNI; or if not available, PVP mesh overlay with required signaling domain between campuses
 - CANs: If ATM and SVC service, fully supported signaling environment and PNNI; or if not available, PVP mesh overlay with required signaling domain between campuses
 5. The IPT will develop the addressing scheme for DON, including a general case (showing all peer groups) and a fleet case (showing all peer groups). The scheme will include a diagram of the PNNI hierarchy. Additionally, the following detailed steps pertain.
 - Obtain an address block from DISA.
 - Support each MAN and all campuses within a region with the address block. Deploy the address block by region.
 - Determine how to allocate address space to the ships and ashore fleet infrastructure.
 6. Establish the MAN routing protocols using a 2-level PNNI hierarchy. Campuses will be able to participate as is appropriate. Show all peer groups for the general case. A similar hierarchy (with peer groups) will be developed for the fleet case. Both will be supported by graphic design guidance and supporting text.

7. Fully define LANE and the cases in which it should be used and not used. In general, LANE will be implemented only on campuses. The IPT will provide design guidance using graphics and supporting text.
8. Fully define MPOA and the cases in which it should be used. The guidance will present the current enterprise architecture for MPOA and provide specific guidance for both the general and fleet cases. It will provide graphic design guidance and supporting text, and diagram MPOA from campus to region to other regions.
9. Determine the additional routing architecture required for the DON enterprise network, including support of autonomous networks and fleet mobility requirements (for ships and for embarked staff, MEUs, air wings, etc.).
10. Develop a DON position for the establishment, configuration, and use of virtual private networks, including the N6 requested strategy and implementation issues (including interoperability, security, management, and performance), and coordinate with the joint services communities to reach a technically sound and implementable solution.
11. Determine a detailed DON connection architecture strategy for voice for both the ATM and IP connectivity environments.
12. Determine a detailed ATM security plan initially focusing on protecting the VPN.

3.8.2 Steps for Developing a MAN

1. Determine governance (including who builds, type of oversight, method of funding).
2. Understand customer base (including geography, population, mission, joint requirements).
3. Analyze requirements (current and projected, media, bandwidth, service delivery points, availability, alternate routes).
4. Develop MAN architecture (including security and management).
5. Identify provider alternatives.
6. Develop test plan.
7. Procure MAN.
8. Establish WAN connectivity.
9. Phase in campuses.
10. Connect outlying sites.

3.8.3 Steps for Developing a CAN

1. Perform inventory of circuits (voice, video, data) and cost of circuits that leave CAN.
2. Document routing architecture (routing domains, autonomous networks, Interior Gateway Protocol, External Gateway Protocol).
3. Determine performance requirements for connection to MAN.
4. Work with the MAN provider to establish Memoranda of Agreement and Service Level Agreements.
5. Determine connectivity strategy to MAN.
6. Determine management of network (if outsourced, turnover visibility and control are included in the MOAs and SLAs.).
7. Align the network topology along geographical versus organizational lines (use MAN to connect CANs in regions, use MAN to connect to WAN service delivery point).
8. Ensure that CAN accreditation is completed.
9. Align router and switching architecture with MAN standards.
10. Clean up local infrastructure to align with the DON ITI architecture.

Volume I, Chapter 4 – Table of Contents

| | |
|---|------------|
| 4. Network Services..... | 4-1 |
| 4.1 Network Time Protocol..... | 4-2 |
| 4.1.1 Service Description | 4-2 |
| 4.1.2 Applicable Standards, Policy, and Guidance..... | 4-3 |
| 4.1.3 Requirements | 4-3 |
| 4.1.4 Assumptions..... | 4-3 |
| 4.1.5 Service Architecture | 4-3 |
| 4.1.6 Roles and Responsibilities | 4-5 |
| 4.2 Domain Name Service | 4-5 |
| 4.2.1 Service Description | 4-5 |
| 4.2.2 Applicable Standards, Policy, and Guidance..... | 4-6 |
| 4.2.3 Requirements | 4-6 |
| 4.2.4 Assumptions..... | 4-6 |
| 4.2.5 Service Architecture | 4-6 |
| 4.2.6 Roles and Responsibilities | 4-8 |
| 4.3 Enterprise Directory Service | 4-8 |
| 4.3.1 Service Description | 4-8 |
| 4.3.2 Applicable Standards, Policy, and Guidance..... | 4-9 |
| 4.3.3 Requirements | 4-10 |
| 4.3.4 Assumptions and Observations | 4-10 |
| 4.3.5 Service Architecture | 4-11 |
| 4.3.6 Roles and Responsibilities | 4-25 |
| 4.4 Electronic Mail..... | 4-25 |
| 4.4.1 Service Description | 4-25 |
| 4.4.2 Applicable Standards, Policy, and Guidance..... | 4-26 |
| 4.4.3 Requirements | 4-27 |
| 4.4.4 Assumptions..... | 4-27 |
| 4.4.5 Service Architecture | 4-28 |
| 4.4.6 Addressing Conventions | 4-29 |
| 4.4.7 Routing Architecture | 4-30 |
| 4.4.8 User Interface | 4-31 |
| 4.4.9 Roles and Responsibilities | 4-32 |
| 4.5 Network News Service using NNTP | 4-32 |
| 4.5.1 Service Description | 4-32 |
| 4.5.2 Applicable Standards, Policy, and Guidance..... | 4-33 |
| 4.5.3 Requirements | 4-33 |
| 4.5.4 Assumptions..... | 4-33 |

| | | |
|--------|---|------|
| 4.5.5 | Service Architecture | 4-33 |
| 4.5.6 | Roles and Responsibilities | 4-35 |
| 4.6 | Web Hosting | 4-35 |
| 4.6.1 | Service Description | 4-35 |
| 4.6.2 | Applicable Standards, Policy, and Guidance..... | 4-36 |
| 4.6.3 | Requirements | 4-36 |
| 4.6.4 | Assumptions..... | 4-38 |
| 4.6.5 | Service Architecture | 4-39 |
| 4.6.6 | Roles and Responsibilities | 4-41 |
| 4.6.7 | Security Guidelines..... | 4-42 |
| 4.7 | File Transfer Protocol | 4-44 |
| 4.7.1 | Service Description | 4-44 |
| 4.7.2 | Applicable Standards, Policy, and Guidance..... | 4-44 |
| 4.7.3 | Requirements | 4-45 |
| 4.7.4 | Assumptions..... | 4-45 |
| 4.7.5 | Service Architecture | 4-45 |
| 4.7.6 | Roles and Responsibilities | 4-47 |
| 4.8 | Public Key Infrastructure (PKI) | 4-47 |
| 4.8.1 | Service Description | 4-47 |
| 4.8.2 | Applicable Standards, Policy, and Guidance..... | 4-47 |
| 4.8.3 | Requirements | 4-48 |
| 4.8.4 | Assumptions..... | 4-48 |
| 4.8.5 | Service Architecture | 4-48 |
| 4.8.6 | Roles and Responsibilities | 4-51 |
| 4.9 | Remote Access | 4-51 |
| 4.9.1 | Service Description | 4-51 |
| 4.9.2 | Applicable Standards, Policy, and Guidance..... | 4-52 |
| 4.9.3 | Requirements | 4-52 |
| 4.9.4 | Assumptions..... | 4-52 |
| 4.9.5 | Service Architecture | 4-52 |
| 4.9.6 | Roles and Responsibilities | 4-56 |
| 4.10 | General Voice..... | 4-56 |
| 4.10.1 | Service Description | 4-56 |
| 4.10.2 | Applicable Standards..... | 4-56 |
| 4.10.3 | Requirements | 4-56 |
| 4.10.4 | Assumptions..... | 4-58 |
| 4.10.5 | Service Architecture | 4-59 |
| 4.11 | Shipboard Voice..... | 4-60 |
| 4.11.1 | Service Description | 4-60 |
| 4.11.2 | Applicable Standards, Policy, and Guidance..... | 4-60 |
| 4.11.3 | Requirements | 4-60 |
| 4.11.4 | Assumptions..... | 4-61 |

| | |
|---|------|
| 4.11.5 Service Architecture | 4-61 |
| 4.11.6 Roles and Responsibilities | 4-62 |
| 4.12Secure Voice | 4-62 |
| 4.12.1 Service Description | 4-62 |
| 4.12.2 Applicable Standards | 4-63 |
| 4.12.3 Requirements | 4-63 |
| 4.12.4 Assumptions | 4-65 |
| 4.12.5 Service Architecture | 4-65 |
| 4.13Multimedia | 4-66 |
| 4.13.1 Service Description | 4-67 |
| 4.13.2 Applicable Standards and References | 4-68 |
| 4.13.3 Requirements | 4-68 |
| 4.13.4 Assumptions | 4-69 |
| 4.13.5 Service Architecture | 4-69 |
| 4.13.6 Roles and Responsibilities | 4-71 |
| 4.14Common Operating Environment Applications | 4-72 |
| 4.14.1 Service Description | 4-72 |
| 4.14.2 Applicable Standards, Policy, and Guidance | 4-73 |
| 4.14.3 Requirements | 4-73 |
| 4.14.4 Assumptions | 4-73 |
| 4.14.5 Service Architecture | 4-73 |
| 4.14.6 Roles and Responsibilities | 4-76 |

This page intentionally left blank.

4. Network Services

The previous chapter provides design guidance for the network connectivity of the DON enterprise network. It includes the physical connections, the protocols that establish connections and maintain communication sessions between systems, and a description of how applications access and inter-operate with the lower-level network communication functions. The Wide Area Connectivity Plan and the Metropolitan Area Network and Campus Area Network templates provide design guidance for network connectivity across the entire Naval enterprise.

The organizations within the DON must receive basic Information Technology Infrastructure (ITI) services in order to support the information requirements pertaining to their basic missions. This chapter describes those basic ITI services that users in all functional areas require that must be accessible from the network on an enterprise basis. These services closely correspond to the Basic Network and Information Distribution Services (BNIDS) described in the DON Information Technology Support Guidance (ITSG). The ITSG is the companion document to this DON ITI architecture and should be used as the DON authoritative source for ITI standards.

All network services must have a common planning framework and consistent implementation strategy. Some services, such as Domain Name Services, must be implemented under an enterprise hierarchical plan. These basic services are described in this chapter with sufficient detail to allow the appropriate ITSCs to consistently plan and implement integrated services. This chapter should be used in concert with the ITSC template in determining ITSC implementation of services.

The network services defined within this chapter are as follows:

- Network Time Protocol (NTP) (Section 4.1)
- Domain Name Service (DNS) (Section 4.2)
- Enterprise Directory (Section 4.3)
- Electronic Mail (Section 4.4)
- Network News / Network News Transfer Protocol (Section 4.5)
- Web Hosting (Section 4.6)
- File Transfer Protocol (FTP) (Section 4.7)
- Public Key Infrastructure (PKI) (Section 4.8)
- Remote Access (Section 4.9)
- General Voice (Section 4.10)
- Shipboard Voice (Section 4.11)
- Secure Voice (Section 4.12)
- Multimedia (Section 4.13)
- Common Operating Environment (COE) (Section 4.14)

4.1 Network Time Protocol

4.1.1 Service Description

Network Time Protocol (NTP) is a protocol that rides on the Internet Protocol (IP) and assures accurate local time-keeping with reference to radio or atomic clocks. This protocol is capable of synchronizing distributed clocks within milliseconds over long time periods. Several DON services require that system clocks are set accurately so that all servers and services have a consistent time. This is especially important for logging servers, file servers, and some security services.

Hosts are assumed to have no other means to verify time other than NTP itself. Although each host typically contains an internal battery-backed clock, a number of factors, including environmental (e.g., temperature), hardware imperfections (e.g., imperfect oscillator), and time setting/resetting inaccuracy cause errors in the reported time. Because there is no synchronization among hosts within the network, local hosts should use internal battery-backed clocks only to confirm the sanity of the NTP time-keeping system, not as the source of the system time.

While some local clocks maintain time-keeping accuracy to a published and trusted standard (“truechimers”), others are consistently slow, consistently fast, or unpredictable (“false tickers”). The accuracy achievable by NTP depends upon the precision of the local clock hardware and stringent control of device and process latencies. NTP provides the means to set and adjust local clocks, thus correcting for offset-slewing, frequency compensation, and other errors in the local clock function. The local clock is then usable by resident applications for the current time/calendar functions, and periodically updates itself based upon previously measured differences between the true time and the local clock time. In this way, a local host continuously “slews” its internal clock to the correct time using calculations from recent NTP updates. Clocks that are relatively stable in frequency need less frequent updates from NTP servers because they can slew to the correct time based upon predictable “offsets” based on error rates. Unreliable clocks “drift” and require more frequent NTP updates to achieve desired clock accuracy.

Clock synchronization over a network requires long periods and multiple comparisons in order to maintain accurate time-keeping. While only a few measurements are usually adequate to reliably determine local time to within a tolerance of 1-2 seconds, periods of many hours and dozens of measurements are required to resolve oscillator skew and maintain local time on the order of a millisecond. Thus the accuracy achieved is directly dependent on the time taken to achieve it. Fortunately, the frequency of measurements can be quite low and almost always non-intrusive to normal network operations. Correctly implemented, local clocks can use NTP to maintain time to within 15 ns and frequency to within 0.3 ms per day.

The DON enterprise time service delivers accurate time via NTP to all servers that wish to subscribe. Even client machines can use this service, if desired. This architecture offers stratum 1 and stratum 2 service at each of the ITSCs. Campus networks are encouraged to install their own stratum 2 or stratum 3 servers for distribution of time service locally.

4.1.2 Applicable Standards, Policy, and Guidance

RFC 1305, Network Time Protocol (Version 3) Specifications, Implementation, and Analysis

A number of applicable references are at <http://tycho.usno.navy.mil>.

ITSG Chapter 6.6.1

RFC1035 – Network Time Protocol (version 3)

4.1.3 Requirements

Must provide a low stratum-level service to any hosts that wish to subscribe to this service. This implies a high degree of scalability, since hopefully all DON enterprise servers will subscribe to this.

Network Time Service must have no single point of failure.

4.1.4 Assumptions

GPS antennas can be installed on the roof at the ITSC site.

4.1.5 Service Architecture

The DON enterprise time service will establish a stratum 1 time source at each ITSC to obtain time from the GPS network. For redundancy, two such time servers will be installed at each ITSC. Because these stratum 1 servers do not necessarily have sufficient performance and scalability to deliver time directly to all subscribers in the region, a hierarchy is established in this guidance.

At an ITSC, stratum 2 servers (a minimum of two) will slave directly to the stratum 1 servers, and to each other, and to NTP sources at one or more ITSCs in other regions. These stratum 2 servers will provide network time to hosts at the ITSC. They can also provide time to any stratum 3 servers in the region.

At the campus level, time servers can be installed if the subscriber population on that campus justifies it. The campus time servers should operate at stratum 2 and slave to the servers in the ITSC. Very small campuses can have their hosts obtain time directly from the servers at the ITSC.

An option currently being explored is to get terrestrial NTP service directly from the U.S. Naval Observatory (USNO). Currently, the network path to USNO is congested and dispersion values over the link are high, thus limiting its usefulness.

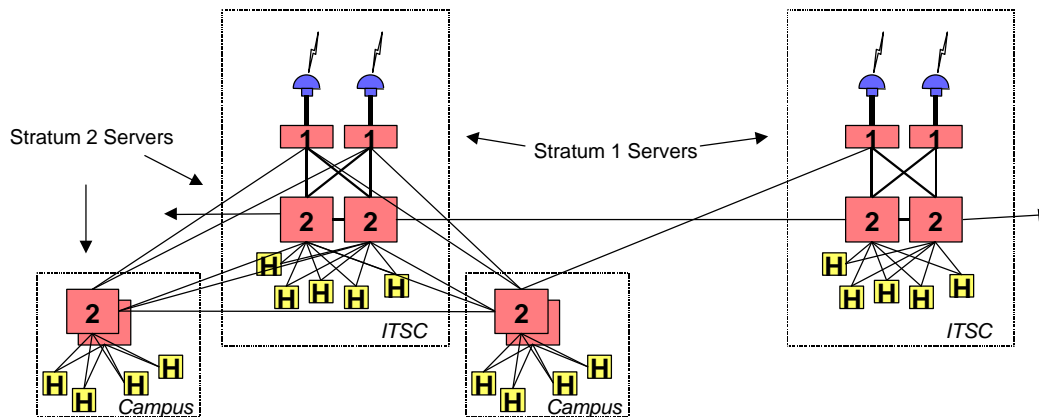


Figure 4-1. Time Service Architecture for ITSC

The architecture diagram in Figure 4-1 illustrates the time service architecture within an ITSC, how campus level stratum 2 servers can peer with the ITSC, and how those servers can deliver time to hosts on the campus.

The DON NTP design is such that accidental or malicious data modification (tampering) or destruction (jamming) of a time server should not result in time keeping errors elsewhere in the DON synchronization subnet. In addition to obtaining time information from multiple, topologically-distributed stratum 1 time servers, all regionally-operated time servers implement access controls within the stratum 2 peer group using the IP security protocol (IPSEC). Redundant timeservers and diverse network paths also increase the quality of the time service within each region. The protocol itself is self-healing because the synchronization hierarchy reconfigures itself after it is determined that a timeserver is no longer in service. Vulnerability is minimized by allowing only designated regional time servers to become synchronization sources for subordinate time servers and denying synchronization requests to/from unknown and untrusted time servers. The NTP standard identifies the available security mechanisms.

The two stratum 1 servers in a region should have DNS names (or aliases) of “tick” and “tock”, i.e., “tick.pacsw.navy.mil”.

Regional Issues and Considerations

Each region will implement redundant stratum 1 and stratum 2 NTP servers as described above. This server will be located inside the DON enterprise firewall.

Regional scalability needs to be considered. Any one server should not be oversubscribed.

Campus and Operational Node Issues and Considerations

Each campus may implement one or more stratum 2 servers using the regional stratum 1 and stratum 2 servers as its primary time source and a neighboring region as the alternate. These time servers should be located on or near the individual subnetworks or LANs containing end systems requiring accurate clocks. For large bases, alternative architectures may be appropriate, including stratum 3 servers on individual LANs within large organizations. Implementation agreements for establishing NTP peer groups for time synchronization are coordinated through the base IT service center. For example, all stratum 3 servers synchronize with each other and lower stratum clocks on the campus and with the regional clocks.

Deployed Forces Issues and Considerations

While in port, deployed forces synchronize with the regional stratum 2 servers just as other operational nodes do. However, because ships are not a full-time participant in the campus time synchronization, they receive synchronization information from other base operational nodes but typically do not send synchronization information to them.

While at sea, deployed forces operational nodes continue to synchronize internal clocks, but terrestrial time sources are unavailable. Instead, ships obtain accurate clock information from GPS or other satellite sources for extended deployments when there is concern that internal synchronization may have unacceptably drifted from actual time.

4.1.6 Roles and Responsibilities

The ITSC will be responsible for providing accurate time within a region. Regional network engineers will be cognizant of the number of time subscribers and ensure that there is adequate hierarchy and redundancy to support the demand.

The ITSC will publish the DNS names of the NTP servers and will provide the instructions for configuring higher stratum servers and for hosts that wish to subscribe to time.

4.2 Domain Name Service

4.2.1 Service Description

DNS is the service that translates domain names to IP addresses and vice versa. A domain name is a mechanism that gives unique names to network devices to eliminate the need for users to remember numerical IP addresses. The DNS service is implemented as a hierarchical distributed database and is accessed using a client/server model. The server component of DNS is the subject of this discussion.

The concept of domain names was introduced at a time when Internet hosts used a flat name space. For example, if one person chose the host name of “eagle,” then no one else on the entire Internet could use that same host name. Clearly, this was not scaleable. Alternatively, the solution was to establish name space domains at the local level where name uniqueness and control only have to be controlled at that level. Domains were assigned to organizations (for example, somecommand.navy.mil). Name uniqueness was guaranteed within that domain because it was controlled within that organization, and uniqueness was guaranteed across the Internet because the name was accompanied by that domain’s name qualifier (for example, eagle.somecommand.navy.mil).

The domain naming scheme is hierarchical. Each level in the hierarchy can assign sub-domains within an assigned name space. For example, the people who control the “navy.mil” domain can allocate subdomains to those Naval commands that have a need for their own domain. This process of establishing sub-domains and assigning control and responsibility to subordinate organizations is called “delegation of authority” and is performed using “NS” resource records in the DNS.

For a given domain, there is always more than one server. One of these domain servers is known as the “primary”, and all others are referred to as “secondary” servers. The primary server contains the master files for the domain (a.k.a. “zone”). The process by which the primary server updates the secondary servers is called “zone transfer”.

4.2.2 Applicable Standards, Policy, and Guidance

See section 6.3 of the ITSG.

See the DISA DNS policy memo for 2nd level domains. (DISA WASHINGTON DC//D//, DTG 162151Z MAR 98)

Applicable RFCs:

- ♦ RFC974 – Mail Routing and the Domain System
- ♦ RFC1034 – Domain Names – Concepts and Facilities
- ♦ RFC1035 – Domain Names – Implementation and Specification
- ♦ RFC1996 – A Mechanism for Prompt Notification of Zone changes (DNS NOTIFY)
- ♦ RFC2065 – Domain Name System Security Extensions

4.2.3 Requirements

In the context of the Naval enterprise, the DNS service must provide the following capabilities:

- Translation of Naval domain names to IP addresses (e.g., hq.navy.mil → 164.224.250.80)
- Translation of Naval IP address space to their respective domain names (e.g., 138.147.50.5 → spider.ncts.navy.mil)
- The above services must be delivered to the entire Internet, not just the Naval enterprise. Therefore, the registered primary servers must have good access to the full Internet.
- Every computer in the Naval enterprise is a DNS client and must associate with one or more DNS servers for general domain lookup services. Therefore, a scaleable architecture is required in order to support 1 million clients.

4.2.4 Assumptions

The assumptions for Domain Name Service will be further refined in future IPT iterations.

4.2.5 Service Architecture

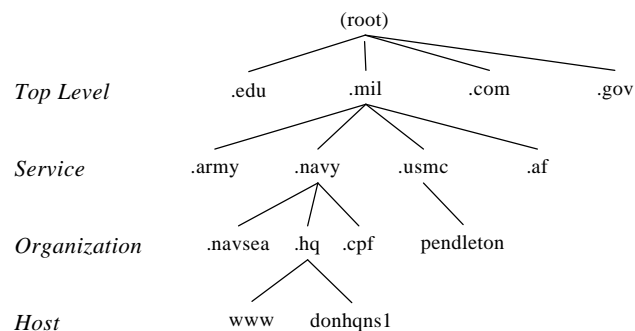


Figure 4-2. Naval Domain Name Hierarchy

The diagram in Figure 4-2 shows the domain naming hierarchy. Organization by service level has been the DoD standard since the mid-1980s. This is strictly a naming hierarchy; it has nothing to do with physical topologies (a common misconception) or distribution of servers.

For each *zone* in the domain name hierarchy (a vertex in the above graph, e.g. `hq.navy.mil`, `pendleton.usmc.mil`, etc.), multiple servers must provide the translation services for that zone. This is a fundamental rule, for reliability purposes, and is enforced by the Network Information Center (NIC). The higher in the naming hierarchy, the more important this is. Therefore, for most organizations, 2 or 3 servers for a zone may be sufficient. But for the second level zones (i.e., `navy.mil` and `usmc.mil`), many more should be established. Note that a single DNS server can support hundreds or thousands of zones. Dedicated servers per zone are not required.

The architecture for deployment of DNS servers will be as follows. Each ITSC will have at least two DNS servers, for both the unclassified and classified (SECRET) levels. These servers will do nothing but DNS. So that these servers can be considered “hardened” for security purposes, all nonessential services (such as send mail) will be disabled and all security patches will be installed. These servers will receive their zone information through zone transfers from the master data source identified for their zone (more about that later). The zone transfers will be protected by DNS Sec. Of the total number of Naval servers (4 at each ITSC), at least 12 should be authoritative (i.e. delegated authority through use of Name Server (NS) resource records in the `.mil` zone). Those 12 (or more) should be chosen for good connectivity and geographic separation. Specific sites will be determined at a future time.

The two primary DNS servers at each ITSC serve as the master servers for the ITSC and the region. All clients in the ITSC will configure their resolvers to point to the master servers. All subordinate DNS servers in a region will be configured to use these primary DNS as forwarders. The primary DNS servers must provide caching services (for performance), as well as recursive queries.

The master data source (per zone) is a machine on which additions and changes are made to the zone information. This is not a critical machine – no clients use it for DNS services. It can be rebooted or restarted at will. It does not need to be a high performance server. The master data source is where the zone administrator makes and tests all zone updates. The primary DNS servers that serve this zone can then obtain the information using DNS Sec-authenticated zone transfers, either in response to an automatic update (based on time-out values in the Start of Authority (SOA) resource record for the zone) or in response to a zone update message from the master data source. This decouples the critical DNS service from the database editing function, and thereby increases overall stability of the DNS.

All of the unclassified primary DNS servers will serve the `navy.mil` and `usmc.mil` zones. Similarly, all of the classified primary DNS servers will serve the `navy.smil.mil` and `usmc.smil.mil` zones. A given ITSC will also service the zones that correspond to commands within their respective region. The authoritative servers for a given organizational zone should include one of the DNS servers at the ITSC, but should also include one or two DNS servers at other ITSCs. In other words, `spawar.navy.mil` might be served by DNS servers in San Diego, Norfolk, and Hawaii. Good geographic separation is achieved in this example.

Naval DNSs outside the DON firewalls will point all queries to the firewalls. This helps to hide the details of the Naval network structure.

Adherence to one rule of naming consistency is required. For every domain name that is returned by the DNS in response to a query to translate an IP address to a name, that name must be a valid domain that can be queried in the DNS, and that query must produce the same IP address as was queried. For example,

if a lookup of the name that goes with IP address 138.147.50.5 returns spider.ncts.navy.mil, then a query of spider.ncts.navy.mil must return the address 138.147.50.5.

Regional Issues and Considerations

Organizations in a region can operate a master data source (primary server). This server must provide accurate data via zone transfers to the secondary server function at the ITSC and must allow local editing of zone information.

Campus and Operational Node Issues and Considerations

A campus may want to install local DNS caching servers to serve on-campus clients.

Deployed Forces Issues and Considerations

The implementation must allow on-board updates, even while disconnected from the network. In other words, the on-board DNS server is the “truth” for that zone. When connected, it provides updates to the NOC or ITSC via zone transfers.

4.2.6 Roles and Responsibilities

The Navy NIC assigns administration of the navy.mil zone. It is currently delegated to UARNOC. For the Marine Corps, it is delegated to the Marine Corps NOC at Quantico.

Administration of organizational zones is assigned to responsible individuals within that organization (if that is the desire of the organization); otherwise it will be provided by the ITSC.

The ITSC is responsible for maintaining the DNS servers (care and feeding of the servers).

The ITSC may be responsible for performing the updates to given organizational zones if that responsibility has been assigned by the respective organization. For example, smallcommand.navy.mil may not have the necessary IT personnel to understand how to run a master DNS, therefore it will obtain this service from the ITSC.

It is the responsibility of anyone connecting devices to the network to properly register IP addresses and names. (ITSCs will registers these devices with the ITCC on behalf of the organizations they service.)

4.3 Enterprise Directory Service

4.3.1 Service Description

A Directory Service provides a function similar to that of a phone book, but generally offers more than just a list of people and their phone numbers. It serves as a repository for “people information” and can be searched to find phone numbers (office, fax, pager, mobile, STU, etc.), e-mail addresses, and mailing addresses. Once the service is in place, it can be used to contain other information attributes about individuals such as passwords, digital certificates, and emergency contact information.

Technologies and standards exist today in support of directory architectures and implementations. Over the years, well-defined protocols and schemas have been established for organizing information in a directory database and for retrieving that information. Modern client applications are becoming “directory aware” and use the standard protocols to locate information. The growing dependence of these client applications on directories is increasing the importance of directories in the technology infrastructure.

Directories can exist in many forms and serve a variety of purposes, but the focus here is primarily on information about people and organizations. The specified DON enterprise directory service will provide a function similar to the “white pages”-- providing the ability to find people in the Navy and Marine Corps by searching on their name. It will also provide the components of a “blue pages” service by providing the ability to “drill down” through an organizational hierarchy and locate groups or individuals based on their organization or billet. A third basic directory function that can be provided is something similar to a “yellow pages” service-- providing information organized by attributes, such as skill, position, mission, or responsibility. The DON implementation of directory service will be called the Naval Enterprise Directory Service (NEDS).

Not addressed here are directory functions such as a directory of file and print services and system configuration information for desktop management that can be found in various Network Operating Systems (NOSs). At present, these directory functions are more of a regional or site issue. When appropriate, these enterprise management functions may be addressed in future versions of this DON architecture document.

4.3.2 Applicable Standards, Policy, and Guidance

- X.500 and LDAP standards – See ITSG section 6.3.2.
- Defense Message System (DMS) standards for Directory Information Tree (DIT) and attributes
- Public Key Infrastructure (PKI) standards for DIT and attributes
- Applicable RFCs:
 - ♦ RFC1777 – Lightweight Directory Access Protocol (v2)
 - ♦ RFC1778 – The String Representation of Standard Attribute Syntaxes
 - ♦ RFC1779 – A String Representation of Distinguished Names
 - ♦ RFC1960 – A String Representation of LDAP Search Filters
 - ♦ RFC2247 – Using Domains in LDAP/X.500 Distinguished Names
 - ♦ RFC2251 – Lightweight Directory Access Protocol (v3)
 - ♦ RFC2252 – Lightweight Directory Access Protocol (v3): Attribute Syntax Definitions
 - ♦ RFC2253 – Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names
 - ♦ RFC2254 – The String Representation of LDAP Search Filters
 - ♦ RFC2255 – The LDAP URL Format
 - ♦ RFC2256 – A Summary of the X.500(96) User Schema for use with LDAPv3

4.3.3 Requirements

This architecture solution must satisfactorily address a mix of enterprise and local directory requirements:

- Must be accessible throughout the entire DON.
- Must be scaleable to support access from hundreds of thousands of clients.
- Should contain entries for all Naval civilian and military personnel and include the ability to add contractors and other associated individuals as appropriate. The resultant number of directory objects is estimated to be approximately 1 million. For each object, the attributes maintained include each individual's full or "common" name, SMTP e-mail address, office phone number, and location (typically city and organization).
- Information in the directory must have a high degree of accuracy based on authoritative sources, and it must be timely (update latency of no more than 24 hours).
- Information in the directory should be searchable based on a person's name. The ability should exist to limit the scope of searches to a major claimant or extend them to the entire DON. The ability must also exist to search by billet, such as "who is the N6 at CINCPACFLT?"
- There should be a natural integration with e-mail. The addressing function of e-mail clients should be able to directly query the directory using standard protocols without the user having to cut-and-paste the results through a separate application.
- The directory solution must support high availability, reliability, survivability, and performance.
- Sensitive information must have controlled access.
- Linkage with the DoD and DON directory hierarchy to the extent possible should support required synchronization of attributes. This requires a directory synchronization function and not a duplicate maintenance of overlapping databases.
- The architecture should enable organizations within the DON to maintain their own directories to meet local requirements, thus accommodating local attributes and control of information.
- Must support a Public Key Infrastructure (PKI) by storing X.509(v3) certificates as attributes to individual directory entries. These certificates will most likely come from the DISA PKI. (The model associated with PKI will most likely grow to hold multiple certificates, such as military identification card and electronic commerce.)
- The directory should be extendable to include new attributes as required by various applications making use of the directory.

4.3.4 Assumptions and Observations

The following important observations and assumptions are relevant to current directory implementations:

- Once an enterprise directory is populated to the extent that it reaches a critical mass, new applications may arise that must be supported by the directory. This dictates that the directory be flexible and have an extendable architecture and implementation. For clarity, "extendable" as it relates to directories is defined to be one that can meet the peculiar directory-related requirements of the Naval or local organization. It also pertains to meeting the special needs of particular mission applications.

- DMS will have a directory solution, but that solution will not meet the DON requirements identified for directories. The DMS directory will be populated only with high-level organizational information in the near term because its primary purpose is to support organizational messaging. In the DON, e-mail services will support the individual sailor and marine, and e-mail must be supported by the directory. There is no requirement for the DON to exactly align with the DMS DIT, although alignment at OU=Navy and above will support synchronization between the DMS directory and the DON enterprise directory.
- DISA is constructing a directory as part of its PKI initiative. While DISA envisions that it will eventually contain all DoD personnel, it will not be able to meet many of the DON-specific requirements listed above. The DISA directory exists primarily to support PKI and organizes directory information under OU=PKI in the DISA DIT.
- There are initiatives in the DON to link selected MS Exchange address books via replication, and this is viewed by some as a potential enterprise solution. However, such a solution is not scaleable to the DON, nor does everyone in the DON use MS Exchange. Also, such replication occurs over proprietary protocols that will be blocked at many firewalls. This replication idea is not a candidate DON solution, however, individual organizations within the DON can continue to perform such replication among themselves to meet local requirements.
- A Navy directory currently exists that is populated with all Naval personnel and receives its data input from the personnel system. Additional populating of directory e-mail addresses and other attributes is often performed by individuals. The result is inaccurate, out of date, or missing information on many individuals. As such, it is not a practical DON solution.

4.3.5 Service Architecture

Because this section is very complex and introduces many new concepts, a number of less-familiar terms are defined below.

- **Naval Enterprise Directory.** The logical architecture and the information used to populate the directory described in this architecture.
- **Naval Enterprise Directory Service (NEDS).** Sometimes shortened to just “enterprise directory.” This is the implementation of the Naval Enterprise Directory. It includes the physical components necessary to provide this directory information as a service to all users. It will be implemented in FY99.
- **Authoritative Source.** This is an accurate source of information for populating the enterprise directory. Many such sources will be employed for implementation of the enterprise directory because any individual source has a limited scope and range of attributes. From the view of the enterprise directory, the claimancy directories (see below) will be used as authoritative sources.
- **Claimancy Directory.** Some of the major claimants in the DON have initiatives underway to provide a claimancy-wide directory or address book. This is often a consolidation of many sources of information from within the claimancy, including various e-mail systems, personnel systems, and NOS directories. These claimancy directories contain most, if not all, of the personnel employed by or associated with a claimancy. Each claimancy should have its own directory. A claimancy directory implementation may exist as a central server or as a set of distributed servers that are replicated from a master, depending on the particular needs of the claimancy. It is strongly recommended that such claimancy directories be “LDAP-enabled” so that information in the claimancy directories can be pulled into the enterprise directory.

- **Command Directory.** Like claimancy directories, some lower echelon commands may have initiatives to synchronize disparate e-mail address books into a consolidated directory. These commands should be feeding their results “up” to their own claimancy directory if it exists, and if not, to the enterprise directory.
- **Regional replica.** A “clone” of the enterprise directory that supports the needs of a given region. These are deployed to the various regions, or Naval concentration areas, as required for performance and scalability.
- **Campus replica.** A “clone” of the enterprise directory that supports the needs of a given campus. These exist mainly for performance reasons, if access to the regional replica does not meet minimum performance criteria. The campus replica may contain only a subset of the objects in the enterprise directory as required by that particular campus.
- **NOS directory.** This is a directory that comes with a network operating system (NOS) for administering things such as user workstations and devices from a network rather than from a host perspective. Examples include NT domain controllers, Novell Directory Service (NDS) in a NetWare environment, and Sun NIS. They often contain people information, in particular, the users of the systems supported by the particular NOS. This can serve as an authoritative source of information for other directories such as command, claimancy, and enterprise.
- **E-mail address book.** Most e-mail systems include an address book function. Often these address books are server-based and contain many names and e-mail addresses of the e-mail system users within a given organization or beyond. This can serve as an authoritative source for e-mail address information for some sets of users.
- **LDAP.** The lightweight directory access protocol. This is the standards-based network protocol used to communicate with the directory. In the client/server model, this is the protocol that the client uses to talk to the directory service.
- **Replication.** Describes the process by which a directory, or a portion of a directory, is “cloned” elsewhere, typically on another server. This allows multiple directories containing copies of some master directory to remain up-to-date. Replicating directories allows scaling to support more users and allows the directory service distribution point to be closer to the user for performance reasons.
- **Synchronization.** Describes the process by which multiple directories containing overlapping information can exchange directory information and remain synchronized.
- **The Directory Information Tree (DIT).** This is the logical hierarchy in which the directory information (objects) is organized.
- **Schema.** Defines the rules, naming conventions, and structure of information in a given directory entry.
- **O=, OU=.** These are X.500 abbreviations used to give names to the levels in a DIT structure. “O” stands for organization. “OU” stands for organizational unit.
- **DISA PKI directory.** This is an instance of an enterprise directory within DoD. It exists for the purpose of supporting the DISA PKI initiative and is where the PKI certificates are stored and from where they should be imported into the Naval Enterprise Directory.
- **DMS directory.** The Defense Message System is a system for performing organization-based message traffic within DoD. DMS has its own directory and is based on X.500 standards.

The directory architecture will be based on a client/server model. Users will operate client applications that obtain directory information by accessing the enterprise directory (server). Two forms of access will

be provided. The “native” directory protocol LDAP (version 3) will be the primary and preferred means for accessing the directory. For situations in which a function cannot be performed using LDAP or in which an LDAP client is not available, a web browser-style interface will be used.

The enterprise directory will be centrally managed with a distributed implementation. The distributed implementation allows increased scalability, reliability, and performance by locating the service close to the end user or client application. “Close” is a relative term and may imply “same hemisphere”, “same region”, or “same campus” as dictated by network paths and other considerations. It must be centrally managed to provide a consistent view across the enterprise.

The distributed components of the enterprise directory will be replicas of a central master directory.

The enterprise directory will leverage efforts of the major claimants by using claimant directories as authoritative sources. Claimants will be encouraged to begin or continue internal directory efforts as appropriate.

4.3.5.1 Logical Architecture

A challenge for the architecture is the poor alignment between the X.500 data structures (schema, Directory Information Tree (DIT)) and the real world requirements. The DIT dictates the boundaries for replication, control, logical structure, search root, and others. In any Naval implementation, the boundaries must be established based on the organizational entities that control the attributes. So, the DIT structure is a compromise that best meets requirements.

The DON directory will use the DIT hierarchy depicted in Figure 4-3.

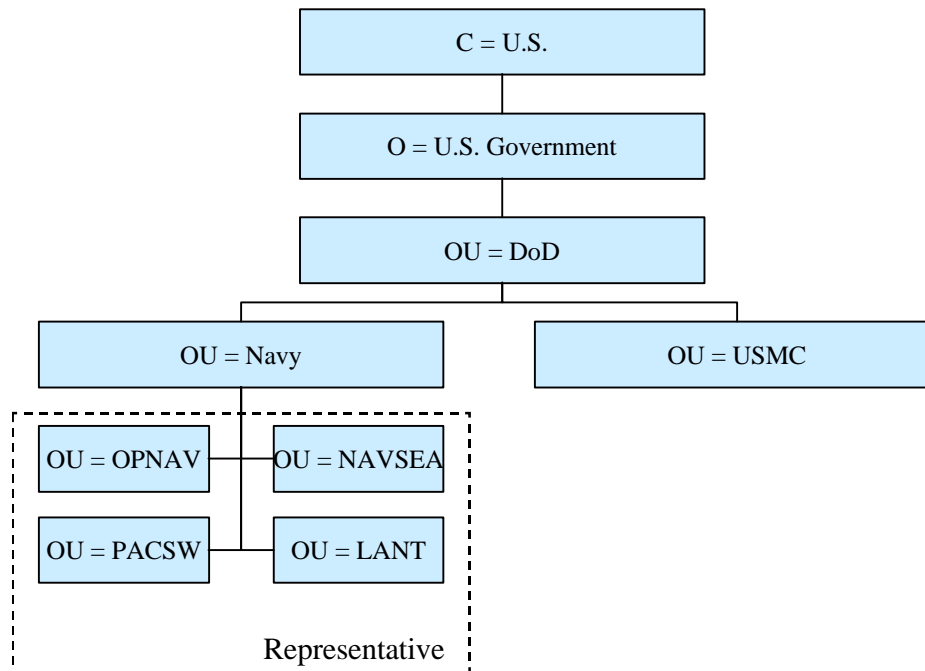


Figure 4-3. DIT Hierarchy

The DIT hierarchy is similar to the DMS structure except that there is no “OU=Organizations” or “OU=Locations” level. The level under “OU=Navy” applies to echelon 1 and 2 commands (major claimants). Below this level, the DIT structure is implemented at the discretion of the particular organization. It is recommended at the lower levels (below major claimant level) that the structure categorize people under “OU=People”, groups of people categorized as “OU=Groups”, and the internal structure of the organization under “OU=Organization”. Other conventions will be established in the future.

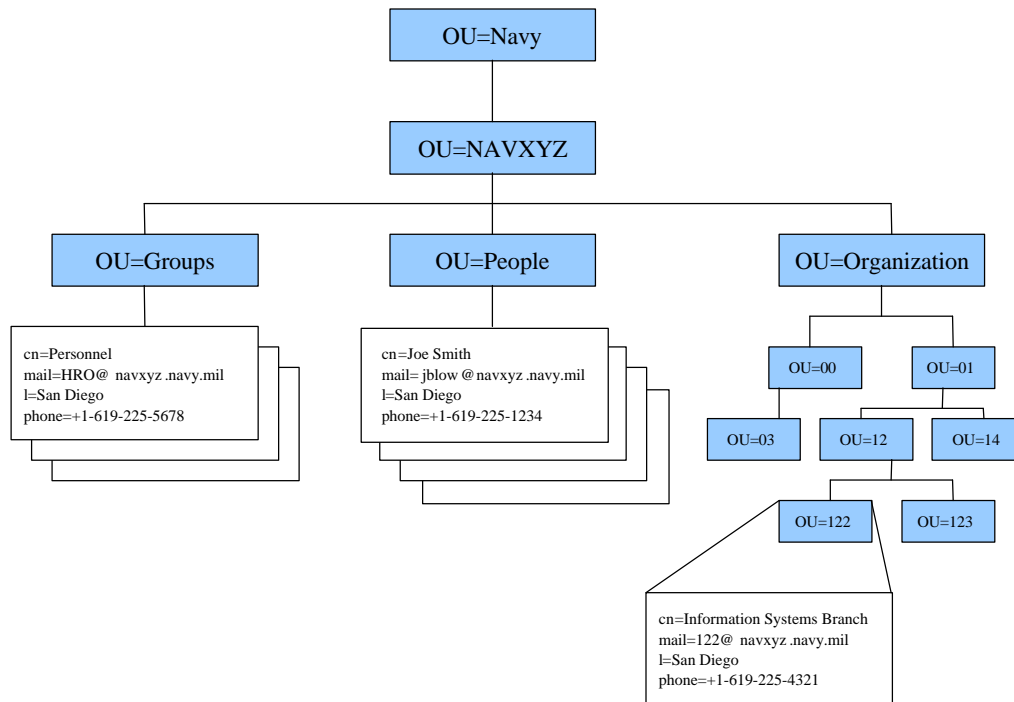


Figure 4-4. Example of a DIT structure below claimancy level

At the same level of the major claimants, there will be OU entries for the various Naval regions such as “OU=PACSW”. This will accommodate regional specific information that does not necessarily need to be replicated globally.

Every object stored in the directory has a name that uniquely identifies it within the directory. The naming scheme is analogous to fully-qualified filenames in a file system in which the name of the file and the whole path must be identified. In a directory this name is called the “Distinguished Name”, or simply “DN”. The DN for any object will include a valid attribute for the object itself plus the entire path within the DIT to this object.

With the structure outlined in Figure 4-4, all “person” objects in this directory will have the same “path” from the root of the hierarchy down to the “People” level. Accordingly, there must be some attribute in each “person” object that identifies that object uniquely at that level. The answer is the common name (CN) attribute, which is created from a combination of the person’s name plus his/her e-mail address. This CN can be used to uniquely identify an entry when combined with the DIT path and thus fully qualify the object name. The format of this special CN attribute is as follows:

```
CN=John Q. Public <jqp@somecommand.navy.mil>
```

The information inside the “<” “>” characters is the user’s official e-mail address. (Note that this is a valid RFC-822 format, and is also the PGP key rings format used for publishing PGP directory keys.)

The fully-qualified DN includes the CN and other attributes that fully qualify the CN relationship in the hierarchy. The fully-qualified DN is illustrated in the following:

```
CN=John Q. Public <jqp@somecommand.navy.mil>, OU=SomeCommand,  
OU=Navy, OU=DoD, OU=U.S. Government, C=US
```

It should be noted that this DN scheme is different from the DMS scheme. The DMS scheme generates a unique identifier for each user and includes that in the CN. The DON scheme proposed does not require a unique identifier.

Within the DON directory naming scheme, the attributes that must be populated for each directory entry are described in Figure 4-5.

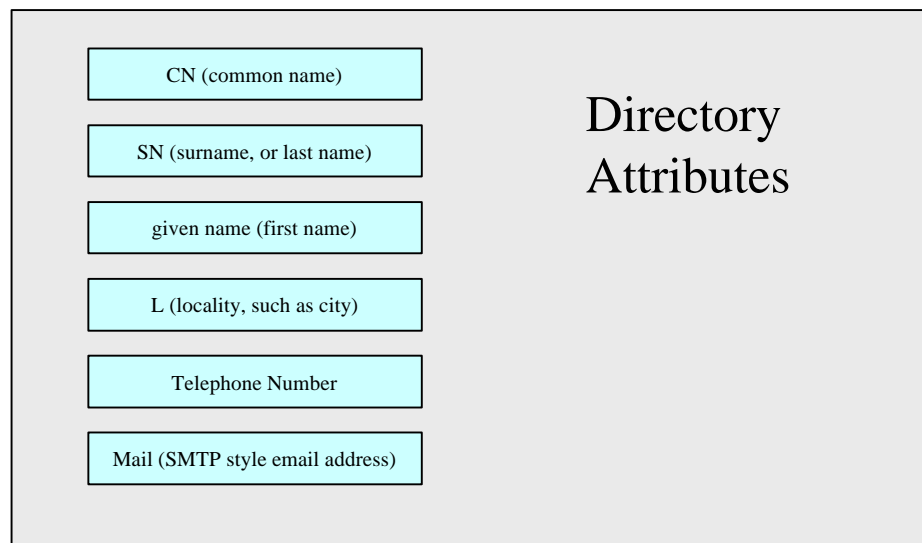


Figure 4-5. Directory Attributes

The CN attribute is multi-valued. Each form of a person's name should be included in the CN attribute to accept queries by all potential searches for that person's name:

CN: William A. Jones

CN: William Jones

CN: Bill Jones

It is envisioned that the major claimants will maintain directories and these will have significant additional attributes to meet local requirements. The ones listed above represent the subset of attributes that are of enterprise-wide interest. The list of enterprise attributes may grow over time as additional enterprise requirements are identified. As additional attributes are incorporated into the enterprise directory, ownership of those attributes needs to be delegated to the authoritative source(s) for those attributes.

If the enterprise directory evolves to the point where unique data is stored in the directory (that is, there is no separate authoritative source), then those attributes will need the appropriate access controls to

facilitate delegation of control to an individual(s) or group that can manage the information in a timely and accurate manner.

4.3.5.2 Physical architecture

The following describes the specific physical implementation of the directory architecture.

- Directory users should be able to point to multiple local redundant directories (local may be on-base or within the region) consistent with user demand and available regional bandwidth. A large campus may want its own directory (or multiple directories) for performance and reliability reasons (similar to multiple DNS servers to provide redundancy).
- A campus replica might include only a sub-tree of the enterprise directory (for example, only the campus's claimancy) because that is the primary information the campus users need to access.
- Each region (fleet concentration area) should have a fully-populated replica of the enterprise directory.
- Replication masters which, in aggregate, fully define a region, should input to the regional directories. There should be multiple masters, such as east coast, west coast, PAC, EUR, and CENT, that replicate down to the 15 regional directories.
- This directory implementation will include a global service that is replicated to the regions. Users can query the regional directory to perform searches and other functions. Major commands that maintain their own directories feed attributes "up" to this global (meta-) directory. The regional directories can replicate only a user-required sub-tree or the whole directory down to a campus, or the regional directory can provide content to a directory administered by a local command or major claimant.

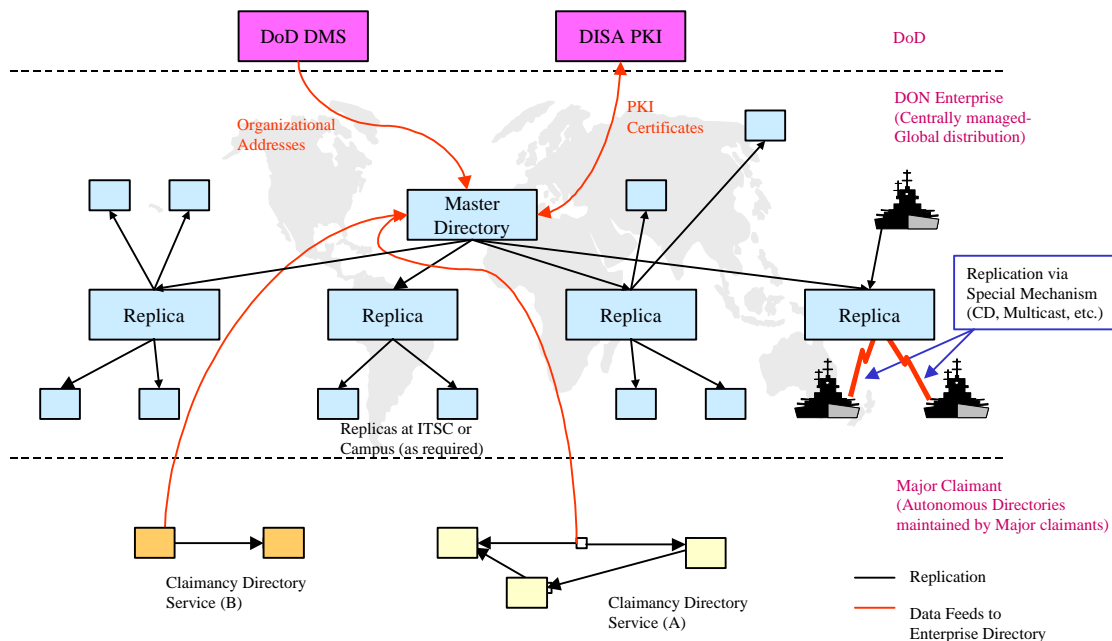


Figure 4-6. Naval Enterprise Directory Physical Architecture

The Naval enterprise directory consists of a three-layer hierarchy. The middle layer as depicted in Figure 4-6 represents the Naval-wide global directory and contains all Naval personnel and associated

individuals. Above the Naval-wide global layer is the rest of DoD. At the layer below are the directories maintained by the major claimants in the DON.

The replication process to the ships will be a modified process which prioritizes the updates by segment and by attribute and makes use of broadcast technologies besides the normal IP connections.

Each major claimant will maintain its own directory (if not now, then certainly in the future). The claimant directories will typically include more information (attributes) than the enterprise level. The loose coupling between the claimant directories and the enterprise directory allows for flexibility and claimancy control. There is a normalization or translation process that feeds the enterprise portions of the claimant directory information up to the enterprise directory. In like manner, the Naval enterprise directory will be coupled to the DoD directory and, with required transformations, provide service input. The DoD directory will import some subset of information from the enterprise directory, and the enterprise directory will import PKI certificates from the DoD directory.

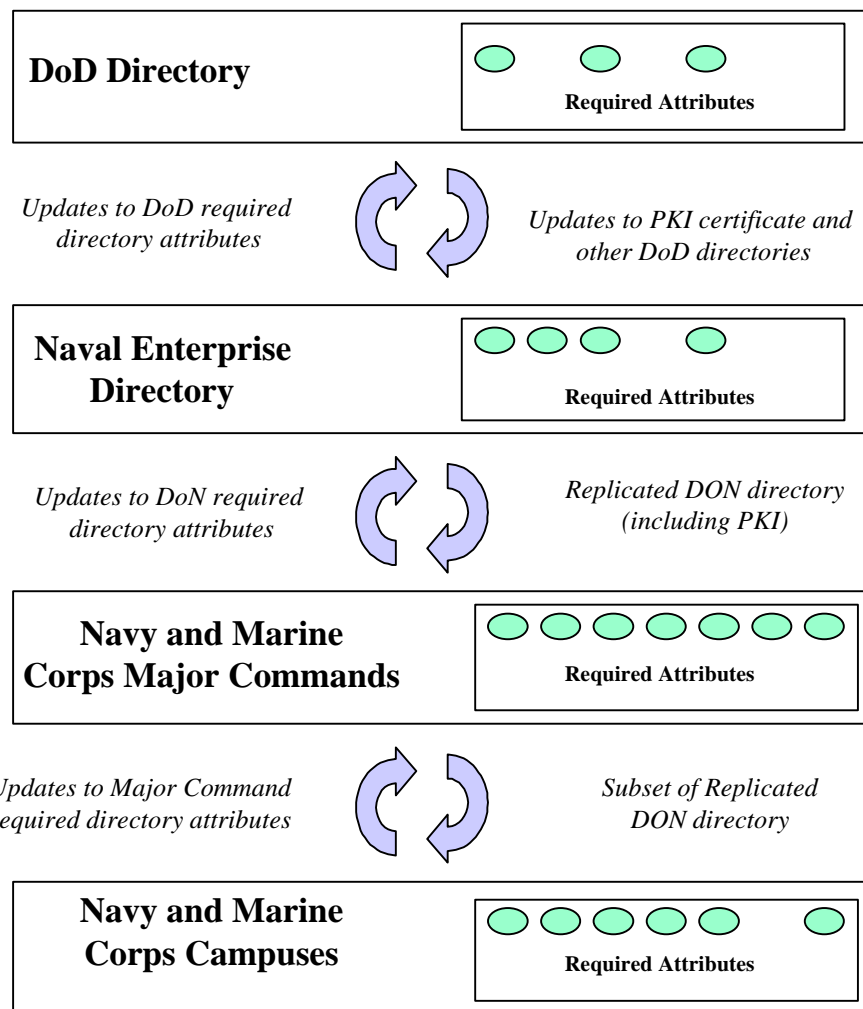


Figure 4-7. Hierarchical Exchange of Directory Attributes

Figure 4-7 shows the replication and normalization of the directory attributes across the multiple information hierarchies. The figure emphasizes the differences in each hierarchy's attribute requirements.

There may be an additional requirement to feed attributes down from the enterprise directory to the claimant directories, but this can be determined on a case-by-case basis because the claimant directory requirements are expected to vary widely.

4.3.5.2.1 Alternative Strategies for Campus Directories

A number of strategies exist for maintaining a campus directory, and the appropriateness of these depends on the individual requirements of the tenant organizations resident on the campus. Figure 4-8 depicts some principal sources of directory information and the following three cases illustrate them.

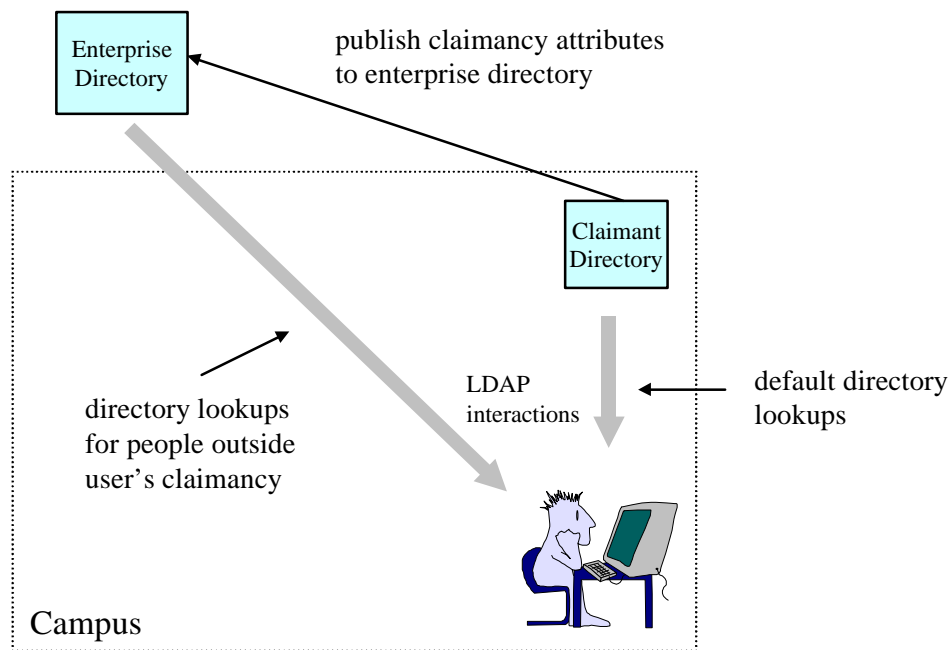


Figure 4-8. Potential Sources of Directory Information

Case 1. The first in Figure 4-8 shows a user that normally gets directory information from his own claimancy directory, which is located on the user's campus (although it could be located almost anywhere). In this case, most directory queries are for information that is within that user's own claimancy.

When the user needs to look for information in other claimancies, he or she goes to the enterprise directory for the information by manually selecting a part of the directory hierarchy that references the enterprise directory, or through a referral process that is implemented in the claimancy directory.

If performance of the enterprise directory is not acceptable, it is possible to obtain an on-campus replica of the enterprise directory (see next example).

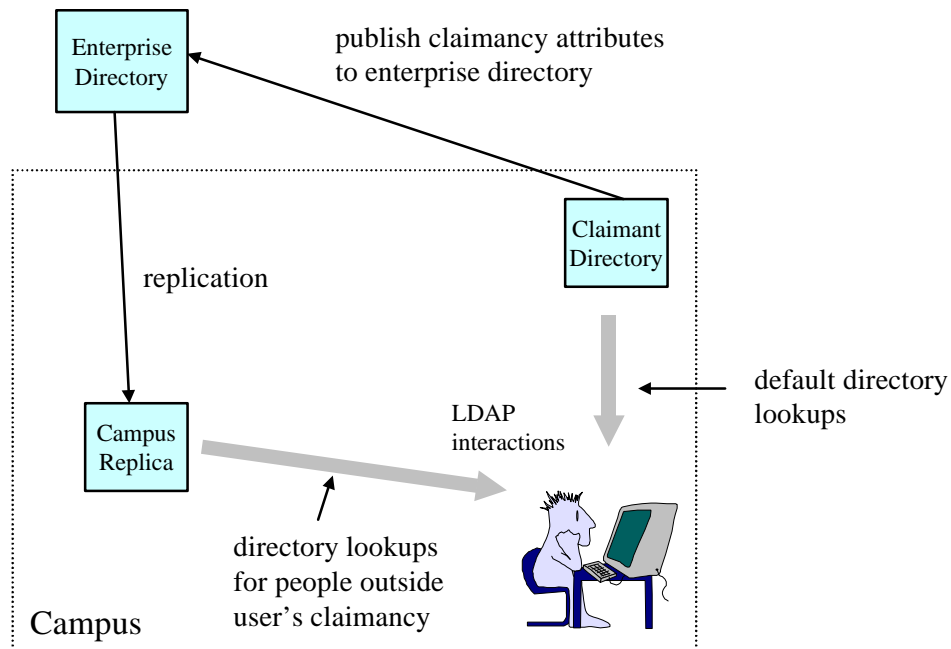


Figure 4-9. Potential Sources for Directory Information

Case 2. This situation is like the previous one except that here there is an on-campus replica of the enterprise directory. It contains the same information as the enterprise directory, but it is physically much closer, and will likely provide better performance.

Because the directory is local, lookups outside the user's claimancy will be performed against this directory.

The idea of merging the two on-campus directories is a possibility if they have consistent DIT structures.

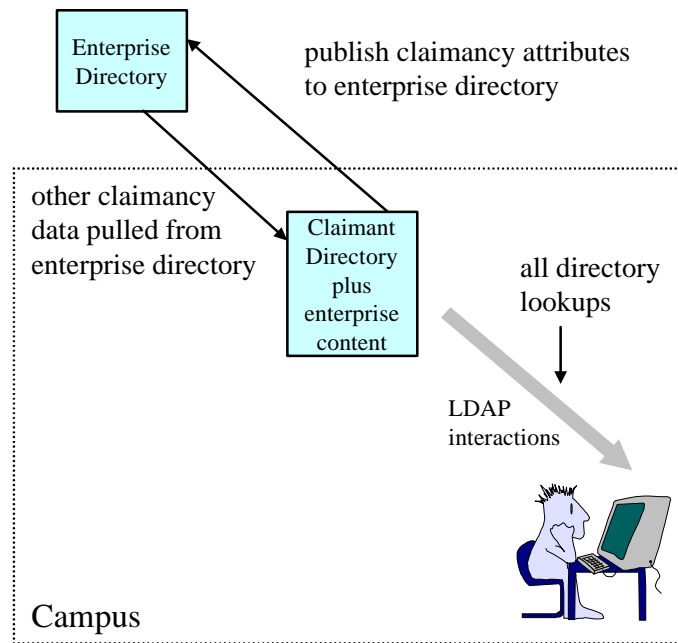


Figure 4-10. Potential Sources for Directory Information

Case 3. This situation in Figure 4-10 is similar to the previous case except the enterprise directory information is now merged into the claimancy directory. The information is “pulled” from the enterprise directory as required. There is no directory “join” process necessary because additional claimancies are added to a directory tree that is already established, assuming the claimancy directory conforms to the enterprise directory structure.

The claimancy still “owns” and manages the claimancy directory. From the enterprise perspective, it is not a true replica because the enterprise is not controlling the replication process, but that poses no problem for the claimancy. The claimancy has achieved a locally-controlled directory that is richly populated and still under claimancy control.

From the user perspective, they can find all DON information in a single directory. This may not be immediately apparent because they still have to navigate the DIT to select the scope of searches.

4.3.5.2.2 Alternatives for Enterprise Directory Replication

The means by which the campus receives replication in Case 3 above depends on the required information and the performance of the network connecting the campus.

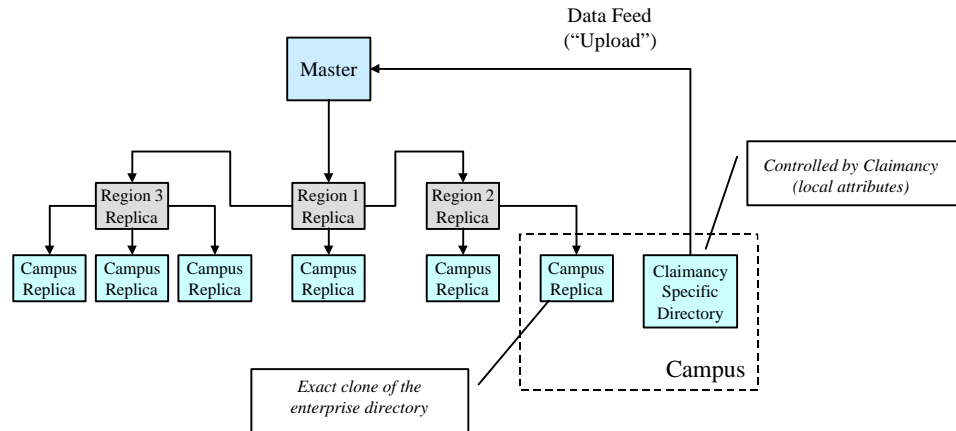


Figure 4-11. Directory Replication

Figure 4-11 shows the data feed from the campus to the enterprise directory. The corresponding feed from the enterprise to the campus should be tailored to the situation. If the updates are strictly controlled, such as a directory data pull situation from shipboard, the specific updates are obtained as required by the campus. Where frequent update is required and bandwidth supports it, a data push from the enterprise directory down to the claimancy or campus directory may be more appropriate. There are accompanying security implications here and these are supported by certificates to ensure integrity of the data.

4.3.5.2.3 Naval Enterprise Directory Service Scalable Architecture

This architecture is very scalable and can grow in phases as demand increases. Initially it could be a single server and a backup for redundancy. Later, replicas could be distributed to other major locations and these replicas could feed other regional or campus replicas. The emphasis must be on making a fully populated Naval Enterprise Directory Service available to all Navy and Marine Corps organizations.

The rationale for locating replicas is based on a number of factors. The first and foremost is to provide reliable access for all users with response time (performance) measured in only a few seconds. Factors that could potentially contribute to poor performance are overloaded servers (support of too many simultaneous clients), congested networks, network outages, low bandwidth links, and large round-trip delays in the network. Factors that will contribute to good server performance are selecting a combination of software and hardware that can support 1 million entries, few-second response times, and a large population (tens of thousands) of users. To ensure that faulty networks do not degrade performance, replicas should be located near the users.

The replication strategy includes providing replication to all 15 or so fleet concentration areas, plus to the many campus networks or remote sites. Considerations for location of replicas include cost, management issues, and possibly political factors.

4.3.5.2.4 Other Areas for Development

For reliability, the DNS name referring to the directory will have multiple “A” records. The DNS will return the names in round-robin fashion if load-balancing is desired or will return “primary first” in the case of a primary and backup server.

As regional ITSCs are established and IT support becomes less command-centric and more region-centric, new directories will be implemented in support of various functions and applications that are region-specific. Considerable thought is required to determine how best to evolve into an integrated directory solution that meets both the enterprise needs and the regional needs. For the near term, this will be done with multiple physical directories, with some loose coupling to the enterprise directory for attributes that should be shared. Designers and implementers are strongly encouraged to follow the enterprise DIT and schema to the extent possible – to facilitate long term migration to an integrated solution. Especially important is the need to organize people information by claimancy rather than by region, because the capability to search by claimancy should be preserved. It is acceptable to put people information under the regional part of the tree, but it must be associated with the proper claimancy.

When it is necessary to look up a name in the enterprise directory, there must be a means to limit the scope of the search. If the user searches for a “Smith” in the total Navy population, there will be too many hits, and it will be difficult to search through the results to find the right person. Alternatively, the major claimant that person belongs to will be known, for example, NAVSEA, and the search can be conducted for all the Smiths located in NAVSEA. In fact, for any search in the directory, it is appropriate to have some initial selector that specifies the claimancy in which the person is assigned. Also, there will normally be an LDAP client default to the user’s own claimancy. In the same way that you have multiple phone books sitting on your shelf, each one for a different organization, you will want a selector in your client interface that allows you to choose which claimancy to search in, or to search all of Navy. This is possible with modern e-mail/LDAP clients as long as the directory is organized in a way that supports this structure.

4.3.5.2.4.1 Blue Pages

The directory architecture includes a description of the DON “blue pages.” The directory should have a representation of the organizational hierarchy of each major claimant. Each node in that hierarchy includes the following:

- Name of that part of the organization
- Organizational code
- Office phone number and fax number
- Head (or acting head) of that organization
- E-mail address of the organization
- List of subordinate organizations

Each node will contain a list of the employees at that level in the organization. A web-based tool will be used to “drill down” into the organization in order to find the office of interest within the organization. (The challenge is to get accurate organizational structures and attributes from the major claimants.) The appropriate DIT and schema for this type of information are contained in Figure 4-3 above. These DITs provide a repository for this information as well as provide an excellent tool for the major claimants to maintain their organizational structure information. (This requires that good tools be developed so that the claimants can easily maintain their portion of the enterprise directory.)

4.3.5.2.4.2 Security Issues

There are security implications that accompany decisions about the information that should be displayed in the directory. A directory database that is rich enough to be useful also contains the same kind of

information that has been the subject of contention on DoD WWW pages. The appropriate directory security policy is not to remove the information but to restrict its access to authorized users and to make portions of the database available only to certain classes of users. For example, individual e-mail addresses should be a widely-distributed application, but conversely, emergency contact information contains non-public items such as home addresses and phone numbers, and should be screened from all external users.

One of the more important directory functions is to make public keys (x.509 certificates) available to everyone. Equally important is providing traceable authenticity -- if a means for spoofing the public keys is exposed, then authenticity is in question. An authenticity trail must go up the directory tree from the point where the public key is manufactured to the core database and back to anyone who needs to use it.

Regional Issues and Considerations

Regional ITSCs will operate a replica of the enterprise directory.

Regional ITSCs will manage the region-specific portions of the directory.

Campus and Operational Node Issues and Considerations

Local commands must input accurate information to their appropriate claimancy directory. This includes an authentic means for communicating x.509 certificate data from units to the ITSC.

Deployed Forces Issues and Considerations

Deployed forces are constrained by limited and sometimes unavailable bandwidth. Directory synchronization across bandwidth-constrained boundaries must be carefully considered and avoided where possible. However, many such units need access to directory information. The following questions emerge: How much of the global directory needs to be immediately available to deployed units? How accurate and timely does the information need to be? To limit loading over the slow communications links, directory updates should be restricted to those objects and attributes that are most important (e-mail address, certificate revocation lists) and should avoid updating those that are less important (phone number, location, etc.). The full directories of regions where a deployed unit expects to operate should be synchronized prior to deployment. Subsequently, only the changes to the directory should be communicated until return to home port.

There are a number of options for distributing updates to the deployed forces. Initial loading of the directory or major updates can be distributed on a CD-ROM delivered to the ship or via a direct connection when the ship is connected at pier side. Another option is to distribute updates via reliable multi-cast over direct broadcast satellite. Yet another option would be for the NCTAMS to prepare updates of limited scope for delivery over RF links at times when bandwidth is available. Any further design discussion is beyond the scope of this document.

4.3.6 Roles and Responsibilities

The ITSC should operate the regional replicas.

Major commands will need to implement and maintain their own directories and feed that information up to the global directory. Note that SPAWAR, NAVSEA, and NAVAIR already operate their own corporate directories.

DON CIO needs to encourage claimants to establish claimancy-wide LDAP capable directories that are highly accurate and provide the necessary “feeds” to the enterprise directory.

A DON level person or group (i.e. “Directory Architect”) should be established to be the arbiter of attribute and DIT conventions. These should be documented at the Navy NIC.

An interface needs to be established with the DoD level directory initiatives so that we can leverage off each other’s efforts. We will be able to offer content to the DoD level directories (they leverage our efforts) and in return will want to import attributes from the DoD level directories such as PKI certificates. This needs to be well thought-out and coordinated.

An overall “Directory Administrator” needs to oversee and administer the operational aspects of the directory.

4.4 Electronic Mail

4.4.1 Service Description

Electronic mail (e-mail) is the basic service for interpersonal and organizational messaging for use throughout the DON. It employs store-and-forward technology that does not provide real-time information exchange but does provide the capability for high-speed communication of small messages or file transfer. It is not a substitute for bulk transmission of large data files to data processing centers or between application servers requiring data replication.

The components of interpersonal e-mail service include the following:

- a user agent for submitting and retrieving messages
- a message store for temporary storage of messages pending delivery or receipt
- a message transfer agent for the reliable transmission through the store-and-forward messaging application network (which uses the DON enterprise telecommunication network)
- an addressing and routing scheme

The architecture that defines the directory components for search and retrieval of recipient e-mail addresses is provided in a separate services section.

All of these components interact to meet the messaging needs of the DON. The service architecture subsection describes the e-mail components, their configuration and interaction, and provides guidance for planning and implementation.

This e-mail service is distinct from the messaging service provided by Defense Messaging System (DMS), whose primary purpose is to process record message traffic. While the e-mail service described here applies to all Naval users, DMS is intended for use by a small segment of the DON with record message traffic requirements. Interoperability between e-mail service and DMS will be addressed in this architecture document so that the systems will be interoperable to the extent required by users. Architecture guidance for DMS is addressed in separate documentation.

4.4.2 Applicable Standards, Policy, and Guidance

- Section 6.2 of the ITSG document
- Defense Message System (DMS) Recommended System Design Architecture (SDA) Document Release 1.1 (Initial)
- Applicable RFCs:
 - ♦ RFC821 – Simple Mail Transfer Protocol
 - ♦ RFC822 – Standard for the format of ARPA Internet text messages
 - ♦ RFC974 – Mail routing and the domain system
 - ♦ RFC1651 – SMTP Service Extensions
 - ♦ RFC1652 – SMTP Service Extension for 8bit-MIME transport
 - ♦ RFC1653 – SMTP Service Extension for Message Size Declaration
 - ♦ RFC1731 – IMAP Authentication Mechanisms
 - ♦ RFC1891 – SMTP Service Extension for Delivery Status Notifications
 - ♦ RFC1892 – The Multi-part/Report Content Type for the Reporting of Mail System Administrative Messages
 - ♦ RFC1893 – Enhanced Mail System Status Codes
 - ♦ RFC1894 – An Extensible Message Format for Delivery Status Notifications
 - ♦ RFC2045 – MIME Part One: Format of Internet Message Bodies
 - ♦ RFC2046 – MIME Part Two: Media Types
 - ♦ RFC2047 – MIME Part Three: Message Header Extensions for Non-ASCII Text
 - ♦ RFC2048 – MIME Part Four: Registration Procedures
 - ♦ RFC2049 – MIME Part Five: MIME Part Five: Conformance Criteria and Examples
 - ♦ RFC2060 – Internet Message Access Protocol (IMAP) – Version 4 rev 1
 - ♦ RFC2110 – MIME E-mail encapsulation of Aggregate Documents, such as HTML (MHTML)
 - ♦ RFC2298 – An Extensible Message Format for Message Disposition Notifications (MDN)
 - ♦ RFC2311 – S/MIME Version 2 Message Specification
 - ♦ RFC2312 – S/MIME Version 2 Certificate Handling
 - ♦ RFC2342 – IMAP4 Name space
 - ♦ RFC2425 – A MIME Content-Type for Directory Information
 - ♦ RFC2426 – vCard MIME Directory Profile

4.4.3 Requirements

The following list of requirements is by no means exhaustive, in that the full list of capabilities and functionality available in today's leading e-mail products is assumed and not repeated here. However, this list highlights specific additional requirements of the Naval implementation which this architecture must support.

The e-mail system must comply with industry standards and remain consistent with those standards as they evolve. Important examples of these include Simple Mail Transfer Protocol (SMTP), Multipurpose Internet Mail Extension (MIME), Secure Multipurpose Internet Mail Extension (S/MIME), Post Office Protocol Version 3 (POP3), Internet Message Access Protocol Version 4 (IMAP4), and Delivery Status Notification (DSN). This is to ensure a high degree of interoperability with e-mail systems external to the DON.

The e-mail system must be compatible with and use the Naval Public Key Infrastructure (PKI) (section 3.8) for S/MIME and other certificate-based security mechanisms.

All user agents must be able to read and write messages in MIME format. Once a message is in MIME format, no component of this enterprise e-mail infrastructure should modify the message body during any part of the transport of that message or as it is stored on any mail store. The goal is to minimize content loss (data or format) through gateways into proprietary environments.

The e-mail infrastructure must support attachments of up to 10 Mbytes in size and be able to reject messages that exceed this size. Bandwidth-constrained environments are exempted from this requirement and will almost certainly need to restrict message sizes to some smaller value. User agents should be able to reject downloading of messages that exceed the user-specified size.

The e-mail system must provide adequate storage for user mailboxes consistent with user requirements. It should include a per-user quota mechanism to automatically manage the available storage space. It is recommended that the quota be set to 50 MB per user.

The e-mail infrastructure should be web-enabled so that users have the capability to get their e-mail with nothing more than a web browser in circumstances where no other client tool is available.

The e-mail infrastructure must filter out unsolicited commercial e-mail ("SPAM" e-mail) as much as possible.

Users should be able to choose an address style that is organization-independent so that the address can be maintained as users undergo reassignment.

4.4.4 Assumptions

- Between DON users in separate organizations, SMTP will be the primary message transport system for interpersonal e-mail.
- Between DON users within the same organization, this architecture is indifferent to the e-mail system used.

- Defense Messaging System (DMS) will be used initially for organizational record message traffic only.
- Interpersonal e-mail does not require the enhanced security mechanisms (high assurance) offered by DMS.
- The e-mail infrastructure described here is a basic utility service and is highly standards-based, as is the case of other basic infrastructure components. It is not the only e-mail service provided within the DON. It is assumed that ITSCs will offer e-mail services that are integrated with groupware environments such as Microsoft Exchange or Lotus Notes for those users that require those environments. The basic e-mail utility described here is available to all DON users, while the integrated or proprietary e-mail environments will be offered only where required.

4.4.5 Service Architecture

The primary architectural components of the e-mail system include a message store (mail server), an SMTP relay (mail transfer agent or MTA), user agents (UA), and a directory providing “white pages” service. It also includes the protocols required to allow interaction among the various components -- SMTP between the MTAs, IMAP4 and POP3 between UAs and the mail server, and Lightweight Directory Access Protocol (LDAP) for accessing the directory. The e-mail service architecture accommodates interpersonal as well as legacy DoD organizational e-mail, which is currently being migrated from AUTODIN to DMS throughout the DoD. Because DoD defined organizational messaging for the services, it is described here only to the extent necessary to clarify the relationship between the DON service architecture for interpersonal e-mail and the DoD architecture for organizational messaging.

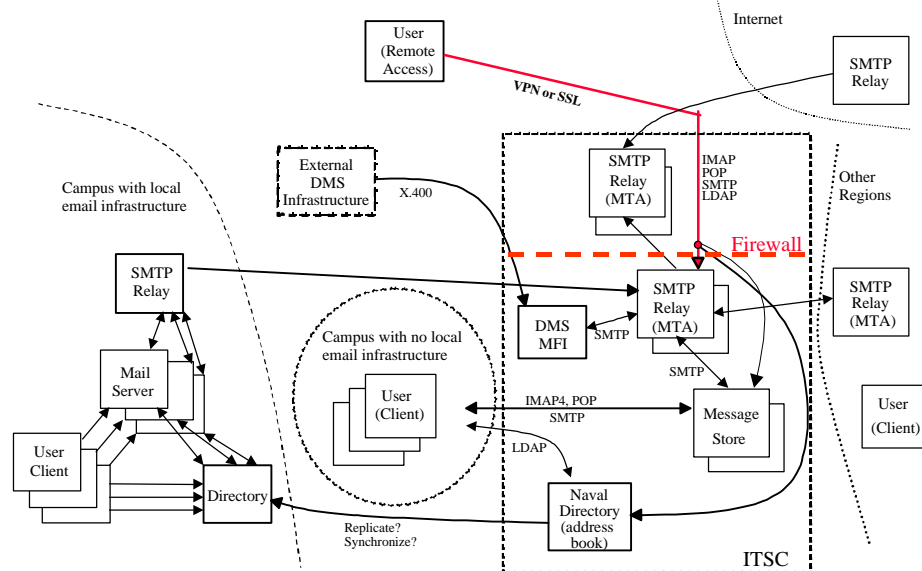


Figure 4-12. Enterprise E-mail Architecture

Figure 4-12 represents the top-level e-mail architecture that shows the interface between the enterprise, the ITSCs, and various user communities. Campuses with and without local e-mail infrastructures are both supported seamlessly. It also shows the interface between the enterprise and external networks. In this regard, network security is provided by various firewalls in accordance with the Defense in Depth strategy.

4.4.6 Addressing Conventions

The addressing conventions used in the e-mail service must first be compatible with RFC-822 and SMTP addressing standards in which the machine-readable address uses the style user@domain. Within the DON, the right-hand side of the “@” sign must be one of the following: navy.mil, usmc.mil, or organization.navy.mil, where “organization” is the name of the organization within navy (e.g. cpf.navy.mil). Sub-domains below the level of organization.navy.mil are possible, but not encouraged. The goal is to keep things flat below the navy.mil level.

For users who require a permanent, unchanging e-mail address despite career/organizational changes (e.g., military personnel), he or she will have an e-mail address in the navy.mil domain. For users who generally stay with a single organization (e.g., civilian personnel), they should have an organization-based address. In both cases, it is imperative that “user” names on the left-hand side of the “@” sign be unique within a given domain.

In the navy.mil domain, the convention for names on the left hand side of the “@” sign will be Firstname.Lastname with the flexibility to allow for conflict resolution or other personal preferences. The rule will be “first come, first served”. This name should not include rank or other designations that are likely to change.

In the organization.navy.mil domains, the conventions for names on the left hand side of the “@” sign will be left to the organizations’ discretion. It is recommended that login names be used in this case because they already need to be unique within that domain. These should be assigned on a first come, first served basis.

All DON e-mail addresses must be registered in the Naval Enterprise Directory (see section 4.3). For each user in the directory, the common name attribute should reflect all names that he or she goes by (e.g. “William F. Smith” and “Bill Smith”) so that directory searches can successfully match all possibilities.

For display names on the “From:” header of actual e-mail messages, the RFC-822 specification allows a number of conventions. Within the DON, the following format will be used:

From: Bill Smith <smithb@clf.navy.mil>

The address inside the “<” “>” characters must be the officially registered e-mail address for the user, and should not contain the mail server name:

(wrong) From: Bill Smith <smithb@mailserver3.clf.navy.mil>

Rank or other titles in the display name are acceptable:

From: Capt James T Kirk <Jim.Kirk@navy.mil>

Characters in the display name (such as commas and periods) that force one to quote the name are discouraged, but will be supported:

(discouraged) From: “Smith, William F.” <smithb@clf.navy.mil>

4.4.7 Routing Architecture

The DON e-mail routing architecture is concerned with proper routing and delivery of messages. Issues for consideration are reliability and/or redundancy, security, relationship to DNS, and the processing of errors.

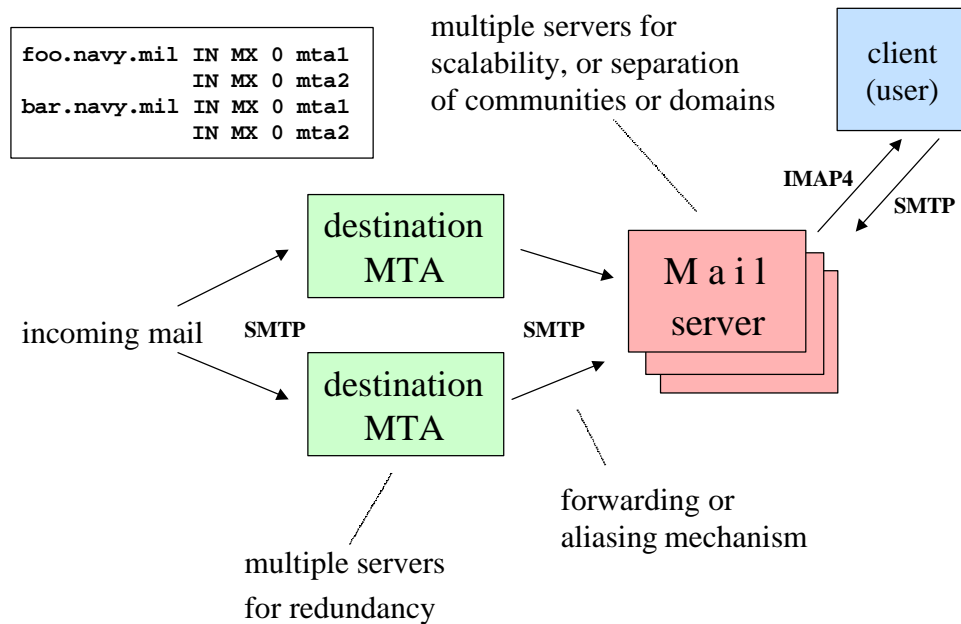


Figure 4-13. E-mail Routing Architecture

For purposes of redundancy and reliability, each e-mail domain will include a minimum of two SMTP servers for receiving e-mail as Figure 4-13 indicates. The DNS will be configured so that all such servers are equally capable of receiving and forwarding inbound e-mail to the destination mail server. Mail Exchanger (MX) records for the e-mail domain must be configured in the DNS and will point at the respective MTAs. The MTAs will be configured so that one can be taken off line with no disruption in service other than a possible degradation in performance.

The destination MTAs and mail servers must be able to support multiple domains simultaneously. The destination MTA forwards the e-mail to the appropriate mail server based on the SMTP destination address of the message. In the process, it hides the internal structure of the e-mail service (i.e. names of mail servers). This allows for a rich and potentially complex internal structure while supporting a simple addressing structure externally. This implies that within a domain, every mail server must know the correct mail server for every user in that domain; accordingly, the server can route e-mail directly to its proper destination without the originator having to specify that in the message headers. A typical solution of this concept is to distribute alias files (mapping to mail server for every user in the domain) on a daily basis to all the mail servers in a domain. A more leading-edge solution is for mail servers to base their routing information on a directory (see direction section) and learn from that central information source which mail server supports a given user. This latter approach is more instantaneous than the alias distribution approach previously described.

A firewall may separate the MTA function into multiple sub-functions. The destination MTA will be configured to disallow relaying between external domains in order to minimize its use in relaying spam.

A special device known as a “Secure Mail Guard”, or SMG, is used to pass e-mail between unclassified and secret networks. They are very restrictive and only allow carefully formatted messages to pass and only between certain source/destination combinations. Attachments generally cannot pass through an SMG. An SMG will be installed at the ITSC in support of this limited capability.

A function that might be implemented in the firewall as part of the e-mail proxy services is a dirty-word search mechanism to allow early detection of classified e-mails leaking to the unclassified side due to human error or lapse of due care.

Users operating an e-mail client will use the IMAP4 protocol to retrieve their mail from the mail server and will use SMTP to send mail via the mail server.

4.4.8 User Interface

The e-mail system must accommodate multiple modes of access to e-mail by individual messaging users. This access can range from direct, on-line access in a LAN or WAN environment to off-line access used when users are unable to establish a network connection (e.g., in flight) and to remote access from the Internet using a web browser on public access workstations (similar to those found in public libraries). The Internet message access protocol (IMAP) defines a client-server standard by which vendors can develop products with the required functionality. Microsoft Exchange, Netscape Messenger, and Lotus Domino are examples of server products that can be configured as IMAP servers. Netscape Communicator is a popular example of an e-mail client that supports IMAP and LDAP.

The e-mail client must be able to work seamlessly with a wide variety of attachment types, including those that are officially registered with Internet Assigned Numbers Authority (IANA) and those that are de facto standards. The user agent must be able to launch the correct application to interpret or display the attachment based on file name or content type. This must include various image types (e.g., gif, jpeg, mpeg, tiff, bmp), document types (e.g., pdf, postscript, rtf), packaging types (e.g., zip, tar), and business application types (e.g., Word, Excel, Powerpoint). At a minimum, plain ASCII, HTML, and RTF must be fully supported in all DON e-mail systems. For other standards, refer to the ITSG.

Regional Issues and Considerations

The mail service provided within a region must scale to accommodate hundreds of domains and hundreds of thousands of users. Domains that are administered within a region must be cognizant of the requirement for uniqueness of user IDs within a given domain.

Campus and Operational Node Issues and Considerations

For any number of reasons, a campus could choose to operate its own mail server(s). It still may choose to use the regional MTAs. The campus network will need to coordinate the mail-forwarding function with the regional provider in this case.

Deployed Forces Issues and Considerations

There are some unique problems with respect to e-mail delivery and routing to deployed forces. The major problem is limited bandwidth. Because of this, restrictions on size of messages need to be applied. Also, message prioritization is a necessary capability to make sure that a large low-priority message does not get in the way of delivering a small urgent message.

Afloat platforms are not always connected. Therefore, a destination MTA must be operated ashore to accept and queue mail for eventual delivery to the afloat platform. The afloat platform should then perform a selective “pull” of queued e-mail when communications are restored.

There are many other issues and details here (in relation to DNS, firewalls, etc.) that are beyond the scope of this document.

4.4.9 Roles and Responsibilities

The DNS administrator is responsible for properly registering the domain and MX records for the MTA host.

The ITSC is responsible for operating a regional e-mail service as outlined above.

One of the ITSCs must be responsible for support of the “navy.mil” domain and should consider using the MTAs of another ITSC for reliability/redundancy of that domain.

E-mail addresses must all be registered in “the directory”.

The directory must provide a “white pages” service to the e-mail clients.

A registration desk must enable users to sign up for e-mail. There is one for each organization and one for the “navy.mil” level.

4.5 Network News Service using NNTP

4.5.1 Service Description

The Network News Service is an information distribution service with which users can selectively gain access to and watch “netnews.” It is different than an e-mail subscription in that the information content is not delivered directly to user mailboxes. That approach does not scale well. With NNTP, all content is stored on a “news server” and replicated to the degree necessary to provide reasonable local access and performance. Client tools “pull” this content from the news servers and make it available to the users for presentation on demand, based on the user’s particular selection of “news groups.” News clients typically can be configured to present only new messages to the user and to organize it by discussion threads within a news group. It also allows contributing to a discussion through “posting” of messages to a news group.

This service is strongly recommended for mass distribution of information or discussion content, especially when the subscriber base is very large or dynamic. NNTP meets the mass distribution requirement much better than the e-mail system because each news article is stored only on a news server and is not duplicated for each recipient.

Discussion groups are organized in a hierarchy. For example, a news group that discusses the NT operating system would be found under “Microsoft systems” under the “computer” category, and would be named comp.sys.microsoft.nt. News groups specific to Navy issues could also be created, e.g. navy.doncio.ipt.itl.

Network news communications between servers and between clients and servers uses the Network News Transport Protocol (NNTP).

4.5.2 Applicable Standards, Policy, and Guidance

ITSG Section 6.8 Network News Transport Protocol

RFC 977 Network News Transfer Protocol

4.5.3 Requirements

Must provide access to the public Internet news groups, of which there are many. May also be required to filter out newsgroups that are not work-related (rec.pictures.dirty).

Must provide the means to create new news groups as required. There must be a means to restrict this capability to a small set of news administrators.

Must provide a “moderator” capability for those news groups that require it.

Must provide the ability to restrict who can “post” articles to a given news group.

Must provide a means to provide authenticated access to certain news groups, if required. Must also provide the means to encrypt information between clients and servers using Secure Sockets Layer (SSL).

4.5.4 Assumptions

It is assumed that DON users will take advantage of this service. Currently, it is severely underused, primarily because people are not in the habit of tuning in to the news service and until recently the performance of news clients for the PC world was poor. With the advent of news support in Microsoft Outlook, Netscape Communicator, and other such systems, it should be easier for users to take advantage of the news service and integrate it into their daily routine.

4.5.5 Service Architecture

The NNTP architecture is quite simple. Each ITSC will have a news server, and these news servers will all be interconnected with NNTP connections for distribution of news. Clients configure their news reader to connect to the nearest news server. Note that mobile users must always connect to the same server, not to the “nearest” one, so that client and server article numbers remain synchronized.

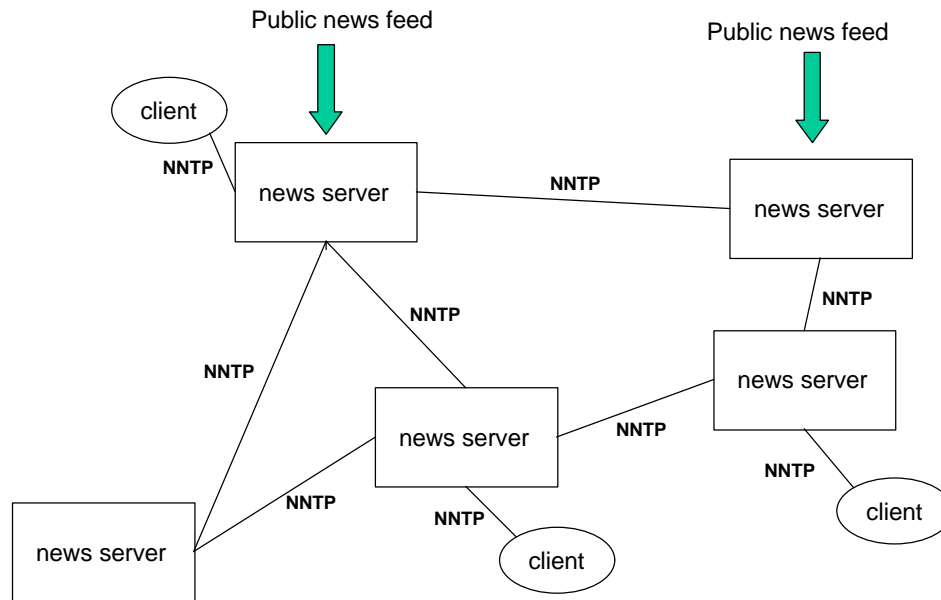


Figure 4-14. NNTP Server Architecture

Figure 4-14 illustrates that no NNTP server has fewer than two peers (to ensure reliability). The architecture includes duplicate public news feeds. One is probably sufficient, but a second is added for reliability or when competing news services offer different content. In each case, all the news will be distributed to all servers.

Most news server implementations provide control over what news groups are sent to peers. This can be useful for regional news groups.

It should be noted that a client can choose to connect to any of the news servers. Also, once a news server is selected, it should remain in use since the client caches information on articles that have been seen, and article numbers on different news servers are not the same.

Regional Issues and Considerations

Each region will have a news server located at the ITSC. The machine should be named “news.domain” where domain is the regional domain name (news.pacsw.navy.mil).

Campus and Operational Node Issues and Considerations

News servers will generally not be located at the campus. Users should have connectivity between their location and the nearest news server.

Deployed Forces Issues and Considerations

The usefulness of the NNTP service to the deployed forces is unclear. NNTP is an interactive, on-demand, high bandwidth service. Certainly while ships are directly connected to the ashore infrastructure (pierside), they have full access to this service like their ashore counterparts. However, while afloat there are serious performance issues because the client/server interface is across the RF communications links. A workable strategy is to host the limited set of news groups of greatest interest on the ship’s news server.

In this way, the news articles are sent just once, and onboard subscribers can pull the information from the onboard news server as required. Where this is appropriate, consideration should be given to deployment of this service to the afloat community.

4.5.6 Roles and Responsibilities

Administrators must be established for management of the news service. These administrators should have the authority to create or delete news groups, to administer access controls to those news groups, and to assign moderators for some news groups if necessary.

Moderators may be required to support some newsgroups.

The ITSC will provide operations support of at least one news server.

4.6 Web Hosting

4.6.1 Service Description

The World Wide Web (WWW) is the basic service for one-to-many information sharing throughout the DON and to the public Internet. It is a “pull” technology that allows individuals to retrieve shared information from servers. The information is stored in Hypertext Markup Language (HTML) or eXtended Markup Language (XML) documents that are moved via the Hypertext Transport Protocol (HTTP) for transport and are displayed across an IP network using a WWW browser.

Running a WWW site involves both the development of the content that will be hosted on the server and administering the web server host and software. Web content development is relatively easy given the COTS what-you-see-is-what-you-get HTML editors. Web site management (controlling the location and structure of the files that make up the web server’s content) is more difficult, but is also facilitated by COTS software. Web server administration (system administration) requires more specialized expertise to keep the web server and its contents secure and in top performance. Many organizations have sufficient hardware and personnel resources to properly maintain their own local web servers. However, for organizations that want a WWW presence without the burden of maintaining a WWW server, the DON Information Technology Service Centers (ITSC) will offer the solution of shared web servers. The ITSC hosted web servers may also be the solution of choice for organizations that want to reduce the risk and resource load of hosting their own public server, but do want to continue hosting their own internal and classified web services. Regardless of the location of the web server, there are basic architectural guidelines (as defined by this document) that should be followed.

Web servers can do more than just provide access to static HTML pages. Web servers can be extended through Common Gateway Interface (CGI) scripts, Java servlets, server side includes (SSI), and other technologies that link application software with the web server. For example, web sites can act as front-end user interfaces for database servers and give easy, forms-based access to data stored in databases to a widely distributed user population. Another common web server function is to serve up a “dynamic” web site. This is a site where most of the content is generated “on the fly” from templates and data is stored in databases based on the user’s specific data requirements. Commercial shopping sites are the most common example of this, but a DON logistics web application or a multilevel intelligence data server can work the same way. This advanced functionality will not initially be available on the ITSC web servers.

ITSCs may provide such specialized services at their discretion. Agreements for these services will be individually negotiated between the ITSCs and their client organizations.

This service is not intended, nor is it a substitute for, bulk transmission of large data files to data processing centers or between application servers requiring replicated data sets.

4.6.2 Applicable Standards, Policy, and Guidance

- Section 6.7 of the ITSG for standards
- DON and local organization ADP security guidelines
- DON and local organization guidance for release of information
- RFC 2068 Hypertext Transfer Protocol 1.1
- RFC 2376 XML Media Types
- RFC 1866 Hypertext Markup Language 2.0
- Internet Engineering Task Force Transport Layer Security (TLS) draft standard (based on Netscape's Secure Sockets Layer protocol)
- Various user authentication, key management and other security standards discussed in other sections of this chapter (see e-mail, directory, and public key infrastructure)

4.6.3 Requirements

No specific web server application or host operating system is specified. However, the web server application should comply with industry standards and remain consistent with those standards as they evolve. Traditionally, there has been little change in web server functionality. The basic web architecture has been flexible enough to support the great evolution in web client functionality without radical changes in the server implementations. Acceptable web server implementations range from "freeware" applications running on personal computers to expensive enterprise server software running on clusters of UNIX systems.

Any organization implementing a web site must comply with the following minimum requirements:

- Must provide a web server that supports the full suite of HTTP version 1.1 services.
- Must have a trained system administrator assigned to maintain the web server host. The system administrator maintains the hardware and operating system and application software on the server and ensures that the system is configured for optimum security.
- Must have a web master or web server administrator assigned to manage the web server software and the contents of the web server. The web server administrator is responsible for implementing proper security procedures as outlined in Section 3.6.7.

- Web servers providing access to the general public must be located outside the organization's firewall and must be on a separate machine from that hosting the organization's private (intranet) web site(s). The intranet web server(s) will be located inside the organization's firewall. Servers outside the firewall should be configured as web servers only to minimize vulnerabilities. The public web server and the firewall software should not be run on the same system. Organizations with high levels of access from other DoD sites may want to consider separating web servers containing that information from the intranet web servers.
- Web servers accessible to the general public should only have Distribution A material (Approved for Public Release) on them. There is also a need to consider the server access controls that apply to safeguarding information with other distribution codes.
- Must provide for access controls on the web site, web site sub-tree, or source directory level based on user name, user domain, or user IP address. Access controls are not required on all sites, but all sites need the ability to implement them. Restricted access controls will be implemented. (There is a problem in that there is not a single consistent mechanism available to all popular web servers to support access controls.) Applicable standards should be applied for implementing encrypted passwords, certificates to increase assurance of user identification, and access control lists (ACLs) for servers that tie to OS access control systems.
- Web server logging functions must be enabled and web server logs must be saved at daily or weekly intervals (depending on site traffic levels) and stored online for at least 30 days. Log files may be stored in compressed format to save space. Historical log files must be saved for at least 12 months. Off line (tape, CD-ROM, etc.) storage of historical logs is acceptable. Web server logs can be useful both in understanding what parts of a web site are being used (to focus future efforts) and to identify suspicious activity.
- Web server log files should be reviewed daily to identify suspicious activity (repeated login attempts, downloading large portions of the web site, ill-formatted URLs, etc.). For small sites with low traffic levels, manual review is possible. Larger sites may benefit from automated log file analysis tools to characterize web traffic patterns and suspicious events.
- The web server application software with all its configuration and support files must be backed up weekly. Web server data files must be backed up daily.
- A web server for public or other organization's consumption must have a robust (reliable and preferably redundant) network connection of sufficient bandwidth to provide reasonable response times to HTTP requests.

In addition, organizations are strongly encouraged to consider meeting the following additional requirements:

- Support certificates or other "strong" authentication means for both users and web servers.
- Support SSL authentication and session encryption.
- Organizations are strongly encouraged to implement change monitoring and/or management systems. These automatically check for unauthorized modifications to the content of the web server and facilitate "rolling back" the server's content to the correct files.

- Organizations should consider implementing integrated text search engines that work across the web site's contents. This functionality is included with the enterprise-level web server solutions and is considered to be essential by many web users.

4.6.4 Assumptions

This guidance is initially applicable to UNCLASSIFIED (NIPRNET) WWW servers. There are some specialized accreditation issues associated with SECRET (SIPRNET) client and server hardware, but the guidance provided here should still be applicable. The SCI level network (JWICS) and WWW system (INTELINK) have their own standards and procedures that are not covered by this document.

Individual organizations must determine the impact of the loss of web servers (internal and external) and the resulting reliability requirements on those servers. DON web servers that contain mission critical logistics or operational data have a high reliability requirement and must have alternate or backup servers in place. Organizational intranets also normally have a need for high availability. High reliability is not typically a requirement for public information servers. High security is a requirement for public servers to prevent embarrassing "hacks" of web sites to introduce false or misleading information. Publicly-accessible servers that are used in the conduct of DON business (contracts and electronic commerce, for example) must consistently be up.

Nothing in this document or the content management process implemented should bypass internal organization approval for publishing information. Access to web servers to manipulate web sites is typically limited to a few individuals who either have standing approval to update elements of a web site or who are the last step in the formal information "publishing" cycle.

A robust DNS exists to allow an organization to create "meaningful" web server names (for example, <http://public.cincpac.navy.mil> or <http://c4isr.spawar.navy.mil>). The DNS also needs to support distribution of an organization's web servers over more than one domain. In a case where an organization's public access server is hosted by an ITSC that is not part of the organization domain (such as the URL <http://public.cincpacfleet.navy.mil>), the DNS might actually take the user to a server located in the disa.mil domain).

A shared directory service will be implemented across the DON infrastructure to simplify the process of maintaining user level access control lists on web sites.

In the case where an organization is hosting web services remotely (at an ITSC), there should be a safe and convenient way to transfer those pages to the ITSC and to have those pages correctly loaded onto the organization's web site. When the author is local to the web server, root, or other privileged log in, this is feasible. This level of access is discouraged for remote hosts. Possible solutions include e-mail or FTP transfer of files to the ITSC for a local web master to load or the use of COTS tools that provide GUI interfaces for remote web site management.

Protecting systems and users from the impact of malicious code delivered via the web is beyond the scope of this initial ITI Architecture. Organizations must be aware of the potential damage that can be done by ActiveX, Java applets, JavaScript, and other executable software that can be delivered as part of web pages. DON policy on these "active" page elements must be developed and appropriate preventative action taken.

Also beyond the scope of this first document are considerations of intelligence and operations security impacts relating to web technology. Potential threats to national and personnel security have already been identified based on the value of information available on unrestricted web sites. Another impact to be

considered is the possible knowledge gained by an opponent through the analysis of individual or organizational web traffic patterns.

4.6.5 Service Architecture

The basic elements of a WWW (or simply web) service are the web client (or browser), the web server, and the local area network (LAN) or wide area network (WAN) that links them. As illustrated in Figure 4-15, the web browser issues a request for a web page from a web server by passing a Uniform Resource Locator (URL) statement onto the LAN or WAN. The routers on the network identify the domain name in the URL and route the URL to the proper web server. The web server parses the URL and places the requested HTML file onto the network addressed to return to the requesting client machine. Typically, web pages are made up of formatted text (the HTML page) and graphics, so the graphics files are also sent to the requesting client where the web browser assembles the web page and displays it. In more complex cases, the URL can instead result in the web server carrying out some action (like a database query) and then returning the results of that action as an HTML page.

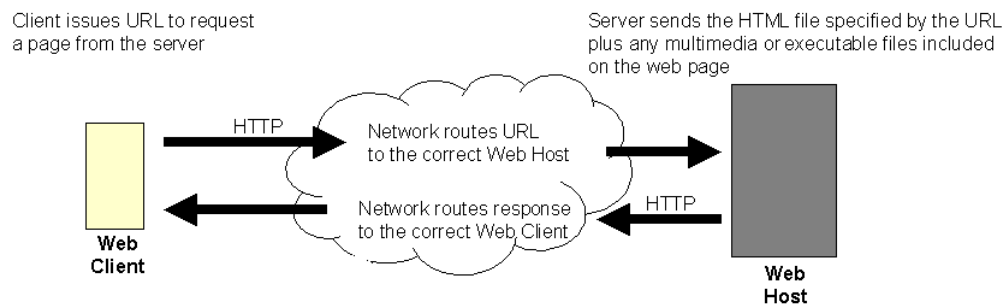


Figure 4-15. WWW Data Flow

For security reasons, most commercial and DoD activities now place web servers that are meant to be accessed by untrusted clients somewhere other than on their internal network. In the case illustrated in Figure 4-16, a firewall has been used to isolate a publicly accessible web server from the organization's internal network that includes a web server for the organization's internal use only. The firewall allows organization members to access web servers on the external wide area network but prevents untrusted users from accessing the organization's internal network and servers.

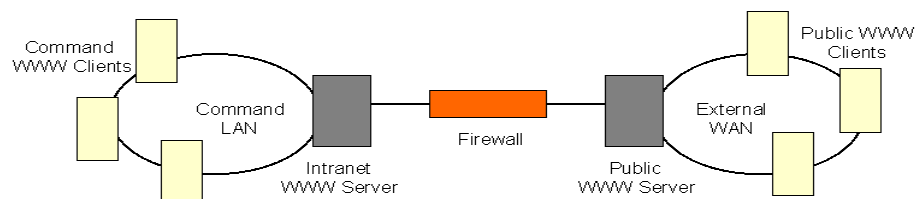


Figure 4-16. Basic WWW Architecture

Figure 4-17 describes a more complex web architecture where a Navy ship has a publicly accessible web site that is primarily used for morale and welfare information and an internal web site for ship's business. Because the morale and welfare site needs to be available when the ship is deployed and because it is designed for access by the general public, the organization has chosen to maintain this web site at a regional ITSC. The ITSC hosts multiple web sites for a number of organizations, so it has a "cluster" of web servers and supports multi-homing. Clustering allows the increased performance to clients on the network by having multiple computers appear as a single web server. Multi-homing allows the ITSC to host multiple web sites, complete with distinctive URLs, on its web server(s).

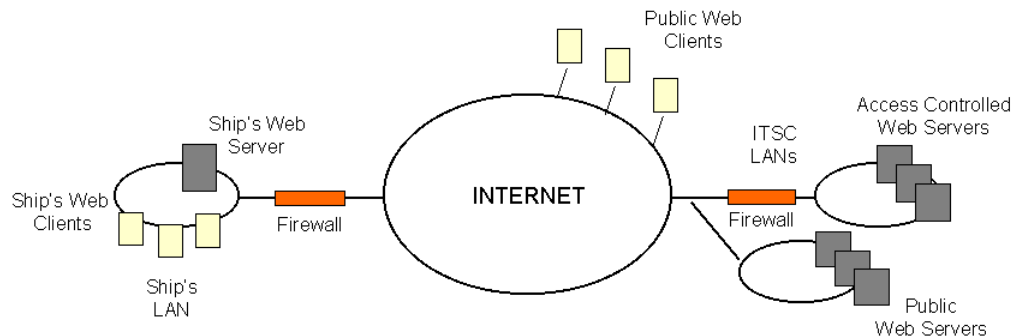


Figure 4-17. Distributed Web Servers

There are three possible web server combinations that an organization can employ:

- A web site hosted by another organization, typically for an organization that only needs a small public web site
- A mix of internal and external web servers - an external organization can host the public access server and the organization can host its own intranet and SIPRNET servers
- Hosting of all of its web services, including multiple servers for public and restricted access

To achieve optimum performance, a web server needs a high bandwidth connection to the network. Enterprise web servers located at ITSCs will all have high bandwidth connections over multiple communications paths. Local organizations may be able to provide sufficiently high bandwidth connections for both their intranet and public web servers.

Other factors that can influence the performance of a web server include the type of network connection, CPU speed, memory, disk space, and the number of other services provided by the server. If a single server is not able to meet client demands (volume of hits or traffic), then a virtual web server can be constructed from multiple server machines. This load-balancing across several web servers can be achieved through proper configuration of the DNS and web server applications. Conversely, several low-traffic web sites (each with their own distinctive domain names) can be combined on a single server by configuring the web server to support "multi-homing."

Web sites with a high reliability requirement should have redundant servers at the local site and a replicated server at a remote site.

Web server functionality is extended through the use of CGIs, servlets, SSIs, and applications that perform dynamic page generation. Each of these options provides increased functionality along with increased security vulnerabilities and complexity. A local organization may be able to host basic web services without any of these server extensions, but a regional facility is almost guaranteed to require them. Even a solution like Microsoft FrontPage, which solves some problems by making web content creation and web site management easier, creates new problems by requiring the FrontPage server extension CGI to be present to achieve all of its functionality.

Organizations should carefully consider implementing any web functionality (for example, Microsoft Active Server Pages) that makes the organization dependent upon one vendor's solution. If an organization insists on using non-standard web functionality and has web sites hosted at an ITSC, the additional costs will have to be borne by the organization.

Regional Issues and Considerations

The web service provided within a region must accommodate multiple domains (multi-homing) and potentially thousands of access-controlled users and tens of thousands of anonymous users. Coordination among organizations to manage domain names and user IDs is essential. Web servers at regional sites must be scalable to provide reasonable response to user requests as the number of domains served increases. The addition of server add-ons (CGIs, servlets, etc.) will require special coordination among the organizations involved and will likely result in additional costs being passed on to hosted organizations.

Campus and Operational Node Issues and Considerations

An organization can choose to operate its own campus web servers, host all its web sites on ITSC (regional) servers, or operate its web sites off a mix of local and remote servers. The variety of security, performance, and quality of service issues addressed elsewhere in this section need to be considered when determining an organization's web strategy.

Deployed Forces Issues and Considerations

It is normally impractical for a deployed force to host a web server due to the constraints of low bandwidth and possibly discontinuous network connections. Deployed forces wishing to make information publicly available should set up a web site at an ITSC. For deployed forces that have a need to access remote DON web servers, caching servers and web subscription services should be employed to make the required information locally available without a need for a real-time connection.

4.6.6 Roles and Responsibilities

Regional

The service provider establishes rigorous technical and administrative controls to ensure that only authorized persons may update published information. Additionally, the service provider makes reasonable attempts to limit access as described above, based upon IP network, domain, and other filtering techniques but cannot guarantee 100 percent success in restricting access. Similarly, it is the information producer's responsibility to abide by security, public affairs, and organization policies for the release of information and to consider the consequences of unwanted release of privacy act or other

sensitive information. The service provider operates and maintains WWW systems, but is not responsible for the information content.

Campus and Operational Node

Campus service providers have the same responsibilities as regional providers.

Deployed Forces Issues and Considerations

Deployed forces will typically not be providing web services. It is the responsibility of the organization to ensure an appropriate web presence is maintained during their deployment through off-board assets.

ITSC

The ITSC is assumed to be the regional provider for web services.

4.6.7 Security Guidelines

Web servers are security vulnerabilities because they are shared resources for “public” access. The security configuration of the web server application must be coordinated with that of the underlying operating system and the network security environment. Weakness in any one of these three areas renders ineffective the precautions taken in either of the other two. The consequences of lax security include loss of data, loss of service, corruption of data, or unauthorized entry into the balance of the organization’s ADP infrastructure through a web site vulnerability.

System and web server administrators must be aware of operating system and web server software vulnerabilities identified by the software vendors and must apply security patches as they are made available.

The following is from the WWW Organization’s Security FAQ available at <http://www.w3c.org> and compiled by Lincoln D. Stein (mailto: lstein@cshl.org).

If you are a webmaster, system administrator, or are otherwise involved with the administration of a network, the single most important step you can take to increase your site's security is to create a written security policy. This security policy should succinctly describe your organization's policies with regard to:

- who is allowed to use the system,
- when they are allowed to use it,
- what they are allowed to do (different groups may be granted different levels of access),
- procedures for granting access to the system,
- procedures for revoking access (e.g. when an employee leaves),
- what constitutes acceptable use of the system,
- remote and local login methods,
- system monitoring procedures, and
- protocols for responding to suspected security breaches.

For Web servers running on UNIX and NT systems, here are some general security precautions to take:

1. Limit the number of login accounts available on the machine. Delete inactive users.
2. Make sure that people with login privileges choose good passwords. The Crack program will help you detect poorly-chosen passwords:

<ftp://ftp.cert.org/pub/tools/crack/>

3. Turn off unused services. For example, if you do not need to run FTP on the web server host, get rid of the FTP software. Likewise for tftp, sendmail, gopher, NIS (network information services) clients, NFS (networked file system), finger, systat, and anything else that might be present. Check the file/etc/inetd.conf (UNIX) or service manager for a list of servers. Deactivate any that you do not use.

4. Remove shells and interpreters that you do not absolutely need. For example, if you do not run any Perl-based CGI scripts, remove the Perl interpreter.

5. Check both the system and web logs regularly for suspicious activity. The programs Tripwire (UNIX) and Internet Security Scanner (UNIX & NT) are helpful for detecting this type of activity:

Tripwire: <ftp://coast.cs.purdue.edu/pub/COAST/Tripwire>

Internet Security Scanner: <http://ciac.llnl.gov/ciac/ToolsUnixSysMon.html>

6. Make sure that permissions are set correctly on system files, to discourage tampering. On UNIX systems, the program COPS is useful for this:

<ftp://ftp.cert.org/pub/tools/cops/>

On Windows NT, consider Midwestern Commerce's Administrator Assistant Toolkit:

<http://www.ntsecurity.com>

7. Consider turning off the automatic directory listings feature of most web servers.
8. Consider turning off the symbolic link following feature of some web servers.
9. Consider turning off the "exec" form of server side includes.
10. Consider not supporting user-maintained directories.

4.7 File Transfer Protocol

4.7.1 Service Description

File transfer protocol (FTP) is used for bulk file upload and download between computers. It is a very simple and efficient protocol that has been in use even longer than e-mail has. From a client perspective, the user can connect to a remote computer and either “get” or “put” one or more files, as well as perform simple file manipulation commands.

E-mail is strictly a “push” technology and has many limitations. The user cannot “get” files at will from a server using e-mail. Also, many e-mail servers limit the size of e-mail messages to a range of 1 MB up to 10 MB. While e-mail is convenient from a sender’s perspective, recipients have little control over receipt of very large attachments that clog up personal mailboxes and degrade the performance of bandwidth-challenged networks.

FTP provides a solution for many of the e-mail shortcomings. Large files can be distributed by placing it onto an FTP server and then announcing a pointer to that location so recipients can download needed files at their convenience.

The FTP service can be URL-enabled. That is, contents of FTP repositories can be referenced with a URL so that easy access is provided through the web browser or similar interfaces.

The FTP service is a streaming protocol. This makes it fast and efficient when compared to other interactive protocols such as a network file service.

In summary, FTP service is useful in those cases in which files are too large to send via e-mail and when the user needs to create a repository of files for users to upload at will.

One important convention in the FTP world is the notion of “anonymous” access. By identifying one’s self to an FTP server as “anonymous” and using an arbitrary password, the user can get public access to FTP repositories. This is very useful when files to be distributed are truly public, and it obviates the need to use passwords and other access controls. Many client tools have this convention embedded as their default authentication, making these client FTP tools very simple to use.

For the DON, the service of FTP repositories is provided for use in distribution of files. It provides both unrestricted (anonymous) and restricted access (limited to authorized authenticated users) to the files in the repository. There is a mechanism for authorized personnel to place files in the repository. There is also a mechanism for unauthenticated (anonymous) users to “upload” to the repository and provide some instructions to the repository administrator for the disposition of such files. There is an administrator of the repository who maintains the overall repository by cleaning out old files, providing the required access to authorized users, monitoring disk usage, maintaining the file structure and indexes, and more.

4.7.2 Applicable Standards, Policy, and Guidance

- ITSG Chapter 6.6.4
- RFC 959 File Transfer Protocol (also known as STD 9).

- http://www.cert.org/ftp/tech_tips/anonymous_ftp_config (How to configure an anonymous FTP server securely)

4.7.3 Requirements

Must align with the service description outlined above.

Must provide both unrestricted (“anonymous”) and restricted access (limited to authorized authenticated users) controls.

Must be “well-connected” to provide high bandwidth access both to the DON network and externally.

Must have an administrator assigned to manage the contents of the repository.

Should follow the conventions used at other Internet FTP repositories in order to allow maximum alignment and utility of COTS products that support these conventions.

4.7.4 Assumptions

High availability is not a requirement, however, a single point of failure probably is not sufficiently reliable.

4.7.5 Service Architecture

The two major architecture considerations are performance and reliability.

To achieve high performance, the FTP server(s) should be well connected to the network. It should be located at an ITSC and be connected at a point that has a high bandwidth external path.

To achieve reliability on an enterprise scale, the FTP service should be replicated at one or more locations. Figure 4-18 depicts this replication process. For general Naval FTP service, there is a primary FTP server at one location, and critical files are replicated on a (very) few servers at other regional ITSCs. For files that are protected through access controls, the access control mechanisms must be synchronized as well between primary and backup servers. Additional work needs to be done to determine the most effective means of automating the dynamic selection of FTP servers, either through Uniform Resource Locators (URLs) listing multiple hosts or through DNS entries with multiple “A” records.

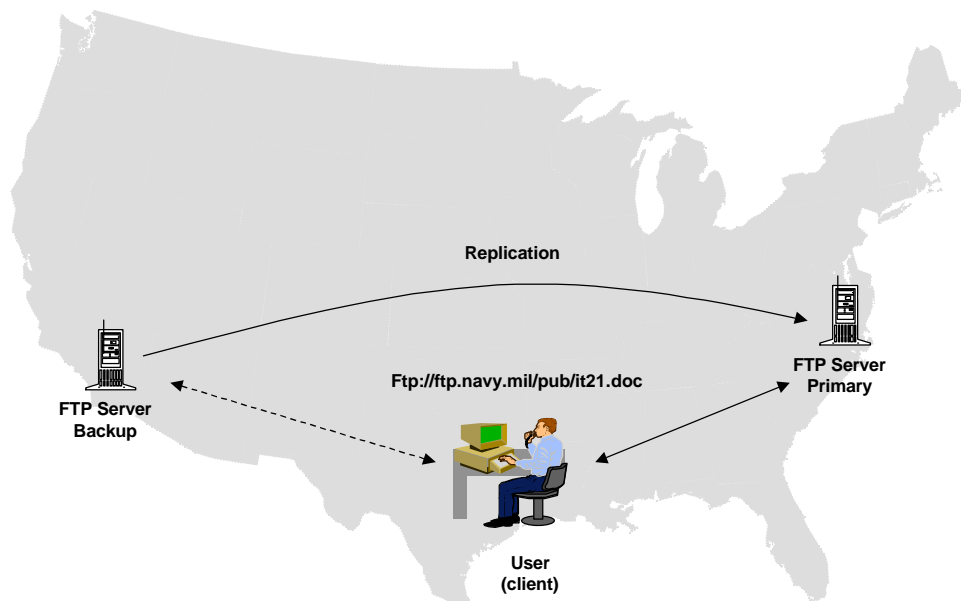


Figure 4-18. Interaction of FTP Components

Regional Issues and Considerations

The FTP service should be provided locally by the regional ITSC. The ITSC operator will need to maintain the server and must conform to enterprise standards for providing this service. For files that need to be distributed strictly within a region, there should be local flexibility for providing this service to the regional customers.

Campus and Operational Node Issues and Considerations

If an FTP service is established within a region, sites that operate existing FTP servers may be able to use the regional service instead. Because a site normally sets up an FTP server for the purpose of distributing files off-site, locating this service at the ITSC is an advantage because it places the service closer to the core of the network. This can off-load the network links between the campus and the rest of the world. Establishing instances of FTP service should be done once per region rather than be duplicated at each campus.

Deployed Forces Issues and Considerations

How deployed forces use this service should take into consideration low bandwidth network links and potentially unreliable service to the deployed platforms. Deployed forces wishing to distribute files via FTP should upload them to one of the FTP servers at an ITSC. Deployed forces that wish to download files via FTP can do so under their own control, depending on available bandwidth. As a quality of service consideration, FTP priority should be lower than interactive traffic when running over IP. This is because FTP may consume all available bandwidth using large packets, which will result in very poor interactive performance (typically small packets). This can be avoided by giving precedence to smaller interactive packets over large streaming packets.

4.7.6 Roles and Responsibilities

The ITSC operator will provide maintenance of the FTP server.

4.8 Public Key Infrastructure (PKI)

4.8.1 Service Description

Public Key Infrastructure (PKI) is the collection of technology, software, hardware, policy, procedures, authorities, and personnel that provides a set of cryptographic tools and the accompanying key management to support digital signature and encryption services to support required applications.

The term “certificate” in the PKI context refers to a data object that binds a public key to a person, a server, or other real-world entity and is cryptographically “signed” by a trusted third party. A certificate supports and extends public key cryptography and can be used to positively identify a person or device and can also be used to implement private channels of communications between individuals.

PKI service provides the mechanism both to generate certificates for individuals and servers and to publish them. There is a “certificate authority” (CA) which manages the life cycle of a certificate. It is the trusted third party that guarantees the authenticity of certificates.

A common directory is maintained to store the certificates for easy retrieval. Additional mechanisms are provided to revoke certificates when required. The CA will publish certificates into this directory.

Certificates are used to provide required services in multiple ways. Typically, they are used for encryption or for digital signatures. In the case of encryption, the PKI service includes “key recovery,” which allows recovery of data in the event of key loss. For digital signature, the service may include the characteristics of non-repudiation. This implies that an individual will have multiple key pairs, depending on the service required.

4.8.2 Applicable Standards, Policy, and Guidance

ITSG Section 3.5 Public Key Infrastructure

PKIX standards:

- See <http://www.ietf.org/html.charters/pkix-charter.html> for a summary of PKIX standards efforts and a complete set of references.

PKCS standards:

- From the RSA FAQ: The Public-Key Cryptography Standards (PKCS) are a set of standards for public-key cryptography developed by RSA Laboratories in cooperation with an informal consortium originally including Apple, Microsoft, Digital Equipment Corporation, Lotus, Sun Microsystems, and the Massachusetts Institute of Technology.
- See <http://www.rsa.com/rsalabs/pubs/PKCS/> for a complete set of references to PKCS standards.

X.509v3 standard format for certificates:

- <http://www.ietf.org/internet-drafts/draft-ietf-pkix-ipki-part1-11.txt>

4.8.3 Requirements

The DON implementation of PKI must meet the requirements of the DoD PKI initiative. The DoD requirements have already been established and are posted on the DoD web site. The basic requirements are summarized as follows:

- Generate certificates in support of Secure Sockets Layer (SSL) encryption for web pages.
- Provide digital signature capability in the e-mail system using the S/MIME protocol.
- Provide encryption capability in the e-mail system (for private or sensitive e-mail).
- Provide positive identification to web servers for access control.
- Provide object signing for Java applets.
- Provide digital signature for electronic forms in support of paperless office initiatives.
- Provide identification certificates in support of VPN access control.
- “Identity certificates” used for digital signature must have the feature of “non-repudiation.”

The DON PKI implementation must publish certificates in a DON directory and these certificates must be retrievable via Lightweight Directory Access Protocol (LDAP). The mechanisms must be established to register users, sign their certificates, and to revoke certificates when necessary.

4.8.4 Assumptions

We assume that it is possible to meet Naval PKI requirements by using the DISA PKI implementation for the DoD. Until recently, this was not considered an option due to the limited scope of the DISA initiative. However, DISA is now committed to meeting Naval requirements and will be incorporated into DON’s plan accordingly.

A DON directory exists and that directory will include all users that intend to use the DON PKI. That directory will be sufficiently reliable, functional, and current to store and retrieve certificates.

4.8.5 Service Architecture

The ITSG describes the components and processes for issuing and revoking certificates. Those details are not repeated here.

The components of the architecture are the Certificate Authority (CA), the Registration Authority (RA), Local Registration Authorities (LRAs), the enterprise directory, and client applications and servers that are PKI aware.

4.8.5.1 Using the DISA PKI

DISA operates the CA hierarchy for DoD. Each of the services has delegated “signing” authority for certificates. This is the “registration authority” function. In the Navy, this function has been delegated to DCMS. Each RA then delegates authority to local commands. This is called “local registration authority”

or LRA. That function will normally reside at the computer help desk, the personnel office, or the pass and decal office for each base.

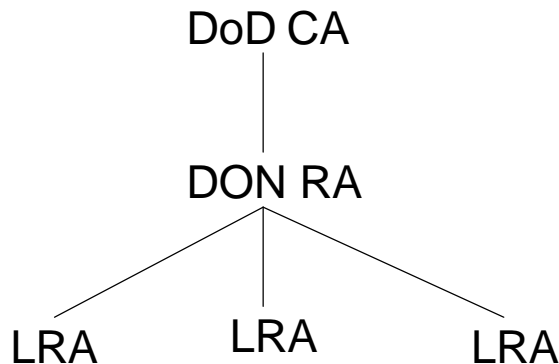


Figure 4-19. Certificate Authority Hierarchy

Figure 4-19 shows the DoD authority hierarchy described above. A similar hierarchy would be employed if the Navy and Marine Corps build their own PKI.

The LRA function will require a significant staffing effort within the DON. It may not require additional personnel, but will certainly require many personnel to take on additional duties. Each Naval location will need to perform the LRA function in support of the personnel at that location. At a minimum they will:

- positively identify each person with multiple picture IDs before issuing a certificate and
- “sign” the person’s public key and perform a process that gets the certificate published in the associated directory.

The exact process is dependent on the particular PKI implementation.

The DISA PKI uses its own directory for publishing certificates. At a minimum, the e-mail certificates generated by DISA must also be published in the Naval Enterprise Directory. This is necessary so that modern e-mail clients that are S/MIME- and LDAP-compatible can send encrypted e-mail messages. In a typical scenario, the e-mail client would retrieve the certificate of each recipient through LDAP access to the enterprise directory or replica. Then the e-mail client can encrypt the e-mail message in the recipient’s public key.

4.8.5.2 Building a Naval PKI

In the event that the DISA PKI does not meet Naval requirements, a Naval implementation will be used. Even though we use the term “implement”, this does not mean that the DON would write all the software from scratch. Rather, COTS products would be used and DON would serve as the integrator of a DON PKI implementation.

With the DON PKI, it is envisioned that multiple PKI implementations will exist. Each of these pilot or prototype implementations will have their own CAs. These are subordinate CAs and will be coordinated and authorized (by signature) by the DON CA. (A possible extension of this arrangement is for an even higher level CA (DoD or National Security Agency) to coordinate and authorize the DON CA.)

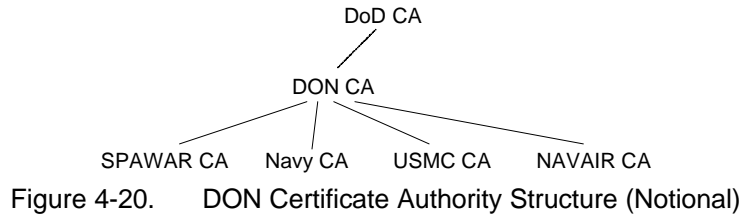
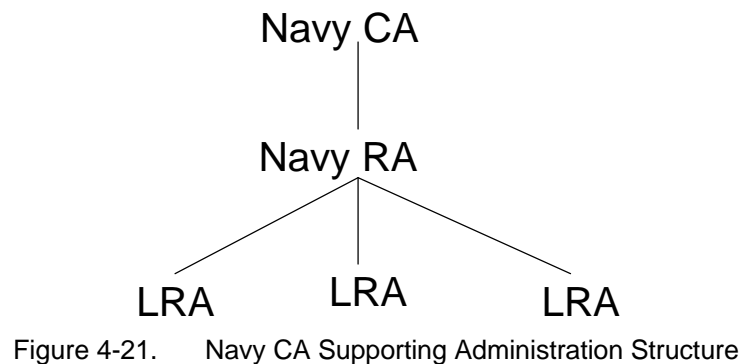


Figure 4-20 is a notional representation of a potential Naval CA hierarchy. The identification of the subordinate CAs requires additional study. What is clear is the need for a defined CA structure.

For each of the eventual subordinate CAs identified, the leaves of this hierarchy include the CA, an RA, and multiple local registration authorities (LRAs). These are depicted for the Navy in Figure 4-21 – those for the Marine Corps would be similar.



The RA delegates responsibility to the LRAs who perform the actual function of signing user and server certificates. The LRA function is performed at locations that normally provide user access, such as a computer help desk or a security office for a base or command.

The directory component is described in detail in section 4.3 of this chapter. The directory contains entries (objects and attributes) for each Naval user, and one of the stored attributes is each user's certificate. This close association of the PKI and directory heavily influences the PKI architecture.

PKI technology is still immature and rapidly evolving. In the DON, there are pilot efforts to gain experience and identify successful strategies for deployment of a PKI. The appropriate tack for this stage of the DON PKI strategy is to maintain a high level view and to appropriately update when required.

Regional Issues and Considerations

The regional ITSC may provide an LRA function.

Campus and Operational Node Issues and Considerations

Each campus will operate one or more LRA functions.

Deployed Forces Issues and Considerations

The deployed forces present a unique challenge for PKI. If user agents obtain certificates from the directory, it implies that these certificates are supported by a directory that is accessible while afloat, or even while no external communications are available. This implies either that certificates must be cached locally or that the entire DON directory must be maintained on-board. Either could be executed – the challenge is to maintain directory synchronization with the master data.

An even greater challenge is the certificate revocation lists (CRLs) that are also stored in the directory. Keeping afloat directories fully synchronized may be difficult due to limited bandwidth, although some studies show that CRL bandwidth requirements may actually be very small and not very dynamic. Solutions must be engineered to ensure that systems which rely on digital signatures properly negotiate revoked certificates, even when there are long delays in distributing CRLs. One example is the case of signed e-mail using S/MIME. If the afloat directory is not getting updated in the correct sequence, the ashore e-mail infrastructure could request signature verification of the afloat unit before the directory updates are received. In that way, messages become invalidated by the ashore infrastructure whether or not the afloat systems stays synchronized. This is only one example. But before we get too concerned about the CRL issue, we must wait and see what the trends are under actual operating conditions. It is the view of some experts that CRLs may actually be very small and not very dynamic, so it may not be a problem.

4.8.6 Roles and Responsibilities

A Naval PKI implementation would require that a DON organization must stand up the DON CA and RA.

PKI implementers must develop a Certificate Practice Statement (CPS) document and ensure that all DON organizations implement it.

Certification Authorities will need to perform policy coordination of CPSs in the event of cross-certification with other organizations.

The RA is responsible for delegating LRAs.

LRAs are responsible for signing certificates for local users.

The directory must allow the CA to store certificates in user objects.

4.9 Remote Access

4.9.1 Service Description

Each region provides a modem pool for telecommuters, travelers, and other users that require dial-up access to the DON enterprise. Regions cooperate to publish local access numbers for all metropolitan areas in all regions. This permits frequent travelers to dial local numbers for access to the enterprise. Secure remote access through the Internet is also provided.

The service provides a directory of local access numbers by region and metropolitan area. A 1-800 number is available for travelers who would otherwise need to call long distance to reach the DON

enterprise network. Virtual Private Network (VPN) service is provided for users connecting via external networks (e.g., those assigned temporarily at a defense contractor site with Internet access and commercial Internet service providers).

4.9.2 Applicable Standards, Policy, and Guidance

RFC 1825, Security Architecture for the Internet Protocol

Interim Guidance for the DOD Public Key Infrastructure, OSD/C3I, 11 Aug. 1998

4.9.3 Requirements

Travelers, telecommuters, and other remote users that lack “local” connectivity need access to the networks and other services described here. Dial-up access must be provided for access through the telephone network. Such dial-up capability must support modern high speed protocols (i.e. v.34, v.90). There must also be a means to access DON services via other Internet Service Providers (ISPs) regardless of whether they are based on dial-up, cable-modem, or wireless access.

There is a need to access “anything from anywhere at any time.” Any authorized individual should have the means to access any DON network or service from anywhere in the world at any time.

These requirements include both unclassified and secret access to DON networks and services.

Each region implements IPSEC-based access using a type-2 security association (see Sec. 4.5 of RFC 1825) between the remote workstation and the security gateway at the boundary between the Internet and the DON Enterprise Network. Regions may implement backup security gateways, but all such implementations shall use the same gateway products and configurations. Implementation agreements between regions ensure consistent configuration and security policies. Remote workstations are configured with IPSEC according to guidelines provided by the regional ITSC.

Authentication is used enterprise-wide so that every region is able to authenticate users throughout the DON. Once travelers obtain access to the DON Enterprise backbone they are able to access resources at their home operational node subject to the Zone 3 and Zone 2 protections, if there are any. Authentication to the DON Enterprise network is equivalent to authentication to their home region, even if the connection is made through another region.

4.9.4 Assumptions

Most travelers will be located in concentrated areas. Therefore, local phone access will be the primary means of accessing the DON Enterprise Network.

All regions implement consistent Zone 4 security perimeter solutions.

4.9.5 Service Architecture

The architecture for providing remote access service consists of the following components:

- Communication Server
- Virtual Private Network (VPN) Gateway

- Authentication Server
- Firewall
- Router

Remote users may access the network through the communication server for dial-up access or through the perimeter router for access from the Internet. In both cases, an authentication server behind the firewall verifies the identity of the remote user. Once the remote user is authenticated, their access to network resources is the same as any other local user except that the firewall may restrict certain functionality (typically based upon protocol, application, or network address) that is deemed unsafe to permit outside the enterprise perimeter. The VPN Gateway is not required in this scenario.

VPN capability is added to the model described to enhance functionality for remote users. The model configuration is shown in Figure 4-22 below. In this case, the remote user establishes a secure tunnel from his remote workstation to his home network as part of the authentication process.

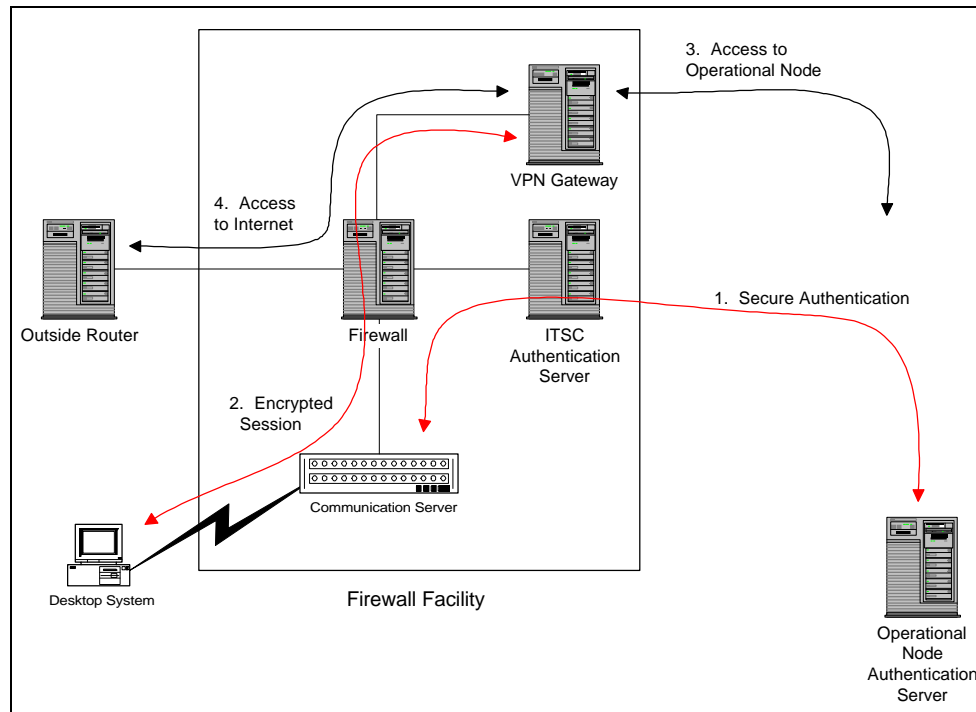


Figure 4-22. Remote Access Service Architecture

First, the user must be authenticated. As shown in Figure 4-22, the communication server establishes a secure channel (shown in red) to an ITSC authentication server. The user provides user ID and password using the remote authentication dial-in user service (RADIUS) protocol or stronger authentication such as SecurID or X.509 security certificate, if required, using a front-end security protocol (RADIUS or TACACS+). If necessary, the ITSC authentication server will employ a back-end security protocol (RADIUS proxy, TACACS+, Kerberos/DCE, Microsoft NT domain security, or Novell NetWare Bindery) to establish authentication within an operational node security context as shown in Figure 4-22. Once the communication server has verified the identity of the remote user, an IP address is issued to the remote user using the DHCP protocol.

Next, an IPSec security association (typically type 2) is established between the remote workstation and the VPN Gateway (shown in red) on his home network, which is located behind the firewall. All traffic (headers and data) between the VPN Gateway and the remote workstation is encrypted and permitted to pass freely through the firewall.

Finally, each end of the tunnel decrypts the received data and forwards it to the correct destination (in the case of the VPN gateway) or application (in the case of the workstation). If the correct destination is outside the firewall, it is again subject to firewall restrictions for outbound traffic (shown in black to indicate no encryption).

For load balancing and redundancy purposes, two (or more) VPN Gateways and Authentication Servers may be implemented.

Communication servers, VPN gateways, and ITSC authentication servers are always located within the regional firewall facility and managed by the regional ITSC staff. For security reasons, communication servers are located outside the firewall. The using community manages its own access control.

Travelers can determine the local phone number for the visited area by dialing a 1-800 number and responding to a menu-driven voice response system. Upon dialing the local phone number, the authentication request is forwarded by the communication server across the DON Enterprise Network to his home authentication server. Upon successful authentication, a secure tunnel is established between the remote user workstation and his home network.

Remote users with Internet service provider (ISP) access to the Internet are also able to establish a VPN connection to their home network through their home VPN Gateway.

To provide Secret access, a similar infrastructure must be put in place, with the addition of approved encryption devices at both ends of each dial-up connection, instead of the VPN solution described above. Secure Data Devices (SDDs) will be installed in-line between the modems and the secret terminal servers. They will need to be permanently keyed to allow for a hands-free auto answer capability. Users will need similar STU-III or SDD devices for secret access and will need these to be registered with the answering SDDs along with MOAs and accreditation paperwork as appropriate to insure proper secure operation and procedures.

Afloat Wireless Access

A special case of remote access is required by ships underway with full term SATCOM or line-of-sight communications. Deployed units which dial in using SATCOM or cellular systems will be treated the same as other dial in remote users are. There are currently four classes of full-term connections that should be supported by fleet teleport facilities. These are distributed as follows:

4.9.5.1 Government SATCOM

SHF - Carriers (CV/N), large deck amphibious ships (LHA/D), Cruisers (CG) (Future)

EHF - All ships except Frigates (FFG)

UHF - All ships

4.9.5.2 Commercial SATCOM

Challenge Athena - Carriers (CV/N) and large deck amphibious ships (LHA/D)

INMARSAT B HSD - All ships

4.9.5.3 Line of sight (LOS) communications

UHF LOS - Amphibious ships and Fast attack submarines (SSN)

DWTS - Amphibious ships

4.9.5.4 Asymmetric communications

GBS - All ships except Frigates (FFG)

Asymmetric submarine communications - Fast attack Submarines (SSN)

Regional Issues and Considerations

Solution planners and designers should seek a single authentication service for dial-up access and Internet access to the DON Enterprise Network.

Regions must establish implementation agreements to ensure consistent (i.e., identical, except for address and other identification parameters) service offerings for remote access via the Internet.

The service architecture described above requires the ITSC to contract with the local telephone company provider to establish sufficient trunks to the ITSC in order to accommodate the region's remote access needs.

Regions maintain an interface to operational node access control databases using the back-end security protocols described above.

Campus and Operational Node Issues and Considerations

Remote workstations accessing the DON Enterprise Network through the Internet implement IPSEC using a bump in the stack (BITS) technique. Regional ITSCs provide configuration guidance.

Access to operational node resources (e.g., file and print service resources) by remote users is facilitated by LAN designs that are IP-based and support access by authenticated external IP addresses.

Campus and operational nodes maintain their own access control databases and provide an interface to the authentication server at the ITSC using the back-end security protocols described above.

Deployed Forces Issues and Considerations

The service architecture components associated with the firewall facility shown above are located in the fleet teleport facility or a shore-based ITSC serving the fleet. Because bandwidth is allocated first for

high-priority operational requirements, authentication and access to the operational node (the ship) may be at a reduced level of performance.

4.9.6 Roles and Responsibilities

Regional

The region maintains all equipment and system components in the firewall facility. It does not maintain individual accounts and access controls but is responsible for providing pass-through authentication to the operational nodes.

Campus and Operational Node

Maintain accounts for secure remote access to local resources. Campus and Operational Nodes do not provide communication servers for remote access. All remote access to the Campus or Operational Node will be come through the regional ITSC firewall facility.

Deployed Forces Issues and Considerations

Allocate bandwidth as needed to maintain high-priority operations.

ITSC

Provide, operate, and maintain modem pools. The total number of modems and associated lines may be reduced significantly as existing capability for remote access is migrated to the ITSC. In order to accomplish this, a rotary system is needed so that all remote users dial a single phone number.

4.10 General Voice

4.10.1 Service Description

Voice is a technology that allows users to communicate interactively in real time through the transmission of sound between two or more users. Voice is a cost-effective tool for enhancing productivity and interpersonal communication between users separated geographically. Typical applications include person-to-person and multi-person real time collaboration, data transmission through the use of modulator-demodulators (modems), and facsimile transmission devices.

4.10.2 Applicable Standards

Information Technology Standards Guidance (ITSG) 98-1.1.

4.10.3 Requirements

The framework for voice will be an open system standard digital switch architecture supporting maximum global and/or regional centralization while ensuring maximum reliability, scalability, and flexibility. Centralization of and resource reductions in administrative functions such as switch management, billing, trouble desk, directory assistance, move-add-change (MAC) technicians, maintenance,

procurement/contract administration, and other voice networking related functions shall be a primary architecture requirement. When possible, the maximum use of existing network resources shall be adopted. However, the elimination of unnecessary or redundant switches, key systems, electronic equipment, or operational/management functions within the existing architecture is also a primary requirement. Sufficient redundancy shall be incorporated to maintain service reliability at or above current industry standards. Additionally, the architecture shall support but not be limited to the following features:

- Common Channel Signaling 7 and/or PRI Trunking
- Multi-Level Precedence Preemption (MLPP)
- Regional E911 Service
- Automated Attendant (If required)
- Centralized Trunking (Public Switch Telephone Network)
- Regional Ashore and World Wide Afloat Number Portability
- Capability of interface within Virtual Private Networks (VPN)

With other switches:

- Voice Mail available for identified users with storage times at current industry standards
- Grade of Service (GOS) at or above current industry standards
- Tail End Hop Off (TEHO)
- Supports Navy BLII initiatives
- Adheres to ITSG 98-1

Within the proposed architecture, regional switches shall be connected as a tandem network (TN) and interconnected via ISDN PRI facilities using switch-to-switch signaling in accordance with open system standards.

The city switch operating within this architecture can be used in a variety of switch exchange applications. It can serve as a local switch which provides end office services, a tandem switch, a toll inter-exchange carrier, an international gateway switch, and/or an Operator Service Position System (OSPS) for national and international calls (with the appropriate software and hardware installed).

The switch must inherently contain the ability, without any further modification, of capturing and displaying emergency on-base calls which shall include the caller's telephone number, the building from which the emergency call originated, as well as the floor and room number of the calling party at a central or consolidated emergency service location.

The city switch's TN feature permits PRI trunks to provide TN trunking between the tandem switch and other switch nodes within the TN. This does not change the implemented TN service for non-PRI trunks, but will partially expand its availability to ISDN PRI users for voice, data, and video traffic. In addition, use of switch-to-switch signaling does not impose any specific topology - the network can be a mesh, star, or main/satellite configuration.

Calls originating on the TN and arriving at the city switch through a PRI will have the following existing features available:

- Automatic Route Selection (ARS)
- Automatic Alternate Routing (AAR)
- Uniform Numbering Plan (UNP)
- Traveling Class Marks (TCM)
- Time-of-day for ARS

Switch-to-switch signaling will provide several identification services between calling and called parties, including:

- Calling Line Identification Presentation (CLIP)
- Connected Line Identification Presentation (COLP)
- Calling/Connected Line Identification Restriction (CLIR)
- Calling Name Identification Presentation (CNIP)
- Connected name identification Presentation (CONP)
- Calling/Connected Name identification Restriction (CNIR)

Other defined switch-to-switch Network Features include:

- Call Completion
- Call Forwarding and Diversion
- Call Interception
- Call Intrusion

The architecture shall also support centralized network management of telephone switches. This management system shall provide but is not limited to the following features:

- Based on COTS software
- Performance management through the collection of statistical data
- Configuration management
- Fault management through detection and isolation of problems
- Security management

The architecture shall also support and provide a mechanism for periodic technology updates based on industry standard timeframes. These periodic technological refreshments shall consider proven technological hardware and software advancements within the telecommunications industry such as voice over IP. However, technology refreshment recommendations must assure continued interoperability within the current architecture.

4.10.4 Assumptions

It is understood that no one business management scenario or overall technological topology will satisfy all the requirements. Due to the diverse mission requirements within the Navy, a case-by-case analysis must be done to determine which management scenario and topology best fits a given region or locality. Whether a regional central office switch, base switch, or Centrex service from the local exchange carrier provides the best solution should be evaluated during the business case analysis process. Additionally, whether a government owned/operated/maintained switch, government owned with contracted operation and maintenance, or completely outsourced service offers the best value to the Navy must be thoroughly examined.

4.10.5 Service Architecture

The proposed architecture features a combination of a city switch and remote campus switches connected together as a Tandem Network (TN). This network solution will provide the Navy with a highly reliable, scaleable, and flexible architecture with centralized network management.

The architecture shall consist of base area networks (BAN), regional metropolitan area networks (MAN), and global backbone networks connected in a hierarchical organization. Additionally, where BAN and MAN are not installed or available, point-to-point tie-line trunking will be used. Each level of the overall network shall support:

- Scalability
- Fault tolerance
- Multi-vendor Open System solutions
- Centralized Network Management

Centralized management shall support the concept of ITSCs as described in Chapter 10 of the ITSC. Each region shall maintain at least one ITSC for central management and administration of the regional network. If required, additional base level centers (ITOC) shall be established to assist the overall regional network management, operation, and maintenance of the regional network. Figure 4-23 provides a basic conceptual overview of the voice architecture.

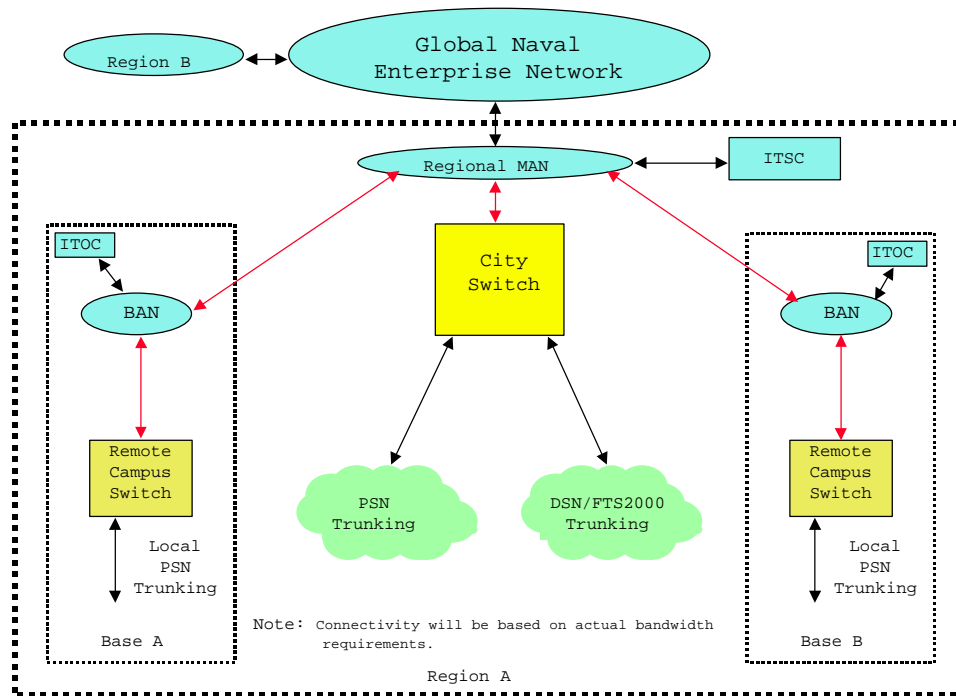


Figure 4-23. Voice Regional Architecture

4.11 Shipboard Voice

4.11.1 Service Description

Shipboard voice communications is the medium that transmits all voice orders and information for intra- and inter-ship communications. Inter-ship communications is generated internally but it also includes the connectivity to off-ship communications.

The boundaries between internal and radio communications are blurring, as are the boundaries between voice, video, and data communications.

4.11.2 Applicable Standards, Policy, and Guidance

Unique military standards for design and procurement have been abandoned to the extent possible to capitalize on commercial investments in systems and infrastructure.

4.11.3 Requirements

Shipboard voice communications is a mixture of mission critical and non-mission critical functions. Mission critical means that the voice capability must survive as long as the ship is afloat. Ships have three states: fighting, steaming, and dead-in-the-water (DIW). Even when DIW, there are intra-ship and inter-ship communications requirements. Industrial-grade products can be integrated such that the result is a highly robust network with a high operational availability. However, the individual user instruments must be robust.

Military unique functions include the following:

- Connectivity and throughput must be sufficient such that for mission critical functions there is no blocking of calls.
- Latency must be short enough that, when a function is keyed with a push-to-talk switch, the first syllable is never missed.
- Preemption is never permitted. Mission critical traffic can originate from anyone.
- Intelligibility must be high enough so that 95 percent of the time, 90 percent of what is spoken is understood. This can be mitigated by speaking in sentences, as is done during telephone conversations in the commercial world, where listeners can extrapolate missing information by using the context of the known conversation. However, on ships, communications are often in phrases or single words. For example, the average call-holding time on a portable communicator is 3 seconds.

4.11.4 Assumptions

Commercial systems and products will be the first choice for solutions and all evidence is that they meet most of the shipboard voice needs. Application of those systems in ships will be such that the mission critical requirement can be satisfied through the robustness of the network. There will be no single points of failure either physically or functionally.

Some requirements can only be met with unique military solutions because there is not a commercial equivalent. One example is that the shipboard voice has a high percentage of traffic in the form of meet-me conferences or nets. There are only a few products that meet that functionality, but solutions can be adapted from available products. The first choice will always be to seek a proven commercial solution that meets the shipboard requirements.

As other facets of communications tend to migrate from independent to systems to network to mission-centric solutions, the same is true of shipboard voice communications. At the highest level, it is the voice function that is needed and specified and not the infrastructure. Today, requirements are often described in terms of solutions, and they should not be.

4.11.5 Service Architecture

The order of priority for defining a satisfactory shipboard voice architecture is affordability, interoperability, and survivability. Within that framework, the reduction in workstation complexity will not only reduce the clutter, but it will also reduce the cost. To achieve that reduction, there needs to be a concurrent increase in the integration of shipboard voice with the other networks.

The goal is anywhere/anytime communications. There is a concurrent migration from wired to wireless communications so that the sailor is always connected to the network and is always available anywhere.

Tactical communications are characterized by communications at short range. Strategic communications tend to be long haul. Mission-critical communications can be tactical or strategic. The same is true of non-mission-critical communications. The need for the connectivity of both will be described in the remainder of this section when completed.

4.11.6 Roles and Responsibilities

NAVSEA will host the Battle Group Engineer for the Navy. That role will be to assure that the focus is on integration of all systems from near term to long term and on requirements versus just solutions.

The top-level focus will be battle management, volume control, and sustained warfighting. These will be decomposed into other functional requirements, but all will have MOEs and metrics. The C4I infrastructure will operate in synergy with the other areas to produce affordable and interoperable battle group operations.

4.12 Secure Voice

4.12.1 Service Description

Secure voice communications for Naval forces ashore and afloat allows geographically-dispersed personnel and activities to securely communicate interactively in real time through the transmission of sound between two or more users. Typical applications include person-to-person and multi-person real time collaboration. Secure Voice equipment additionally supports ad hoc data transfer applications over circuit-switched connections and RF media.

4.12.2 Applicable Standards

| | |
|--------------|---|
| STE-210 | Secure Terminal Equipment Signaling Plan- Interoperable Modes |
| SVS-210 | Signaling Plan-Interoperable Modes |
| FSVS-211 | Interface Control Document for STU-III Black Digital Interface |
| FNBDT | Future Narrowband Digital Voice Terminal |
| FED-STD-1015 | Analog to Digital Conversion of Voice by 2400 Bits/Second Linear Predictive Coding (LPC-10E) |
| FED-STD-1016 | Analog to Digital Conversion of Radio Voice by 4800 Bits/ Second Code Excited Linear Prediction (CELP) |
| CCITT G.721 | 32 kbits Adaptive Differential Pulse Code Modulation (ADPCM) |
| CCITT Q.931 | Digital Subscriber Signal Subscriber No. 1 (DSS1) Network Layer, User-Network Management |
| ITU-T Q.921 | Digital Subscriber Signal Subscriber No. 1 (DSS1) Data Link Layer SR-NWT-001937, Issue 1 National ISDN-1 SR-NWT-002120, Issue 1 National ISDN-2 |

4.12.3 Requirements

Secure voice capability is required in all enclaves; the specific security mechanism employed is dependent upon both the communications capability and mission tasking of the customer. Interoperability among dissimilar networks is typically achieved via inter-working function devices and/or specifically engineered gateways. All systems employed by DON shall be supported in the Global Secure Voice System (GSVS) architecture. Employment by Allied/Coalition partners is subject to technology releasability considerations. Specific secure voice cryptographic equipment systems employed and/or planned for use by the DON include:

4.12.3.1 Secure Terminal Equipment (STE)

Secure Terminal Equipment (STE) is the next generation Secure Telephony device for the U.S. Government. Transition from the current STU-III system is scheduled over the next five to seven years. The STE product line will incorporate four distinct secure telephony modes which each have unique connectivity requirements.

(1) **Secure Terminal Equipment Mode:** The STE mode is a voice service implemented on an ISDN data channel and requires an end-to-end Unrestricted Digital Interface (UDI) (which is 56-64 kbps). In order to use the STE mode, the telephone instrument currently requires an ISDN S/T interface. To support maximum functionality, the provisioning of ISDN service should conform to NI-1/2 industry standards. This mode can be extended to deployed forces via STE direct dial, which is in development.

Early testing of the STE has revealed that government telephone service contracts typically contain no provision for NI-1/2-compliant ISDN. Testing has also revealed that FTS-2000 voice trunking services do not consistently support end-to-end UDI.

(2) **Secure Telephone Unit (3rd Generation) (STU-III) Mode:** The STU-III mode is a voice service implemented on standard PSTN/DSN voice channels. Compression techniques employed on the voice channels must be engineered to support LPC-10 and CELP algorithms. It is fully interoperable with the current STU-III series of equipment and extended to deployed forces via various direct dial methods. It is important to note that the STU-III mode is inferior in quality to the STE mode and that if proper ISDN provisioning is not achieved, the unit will default to the STU-III mode of operation.

(3) **Future Narrowband Digital Voice Terminal (FNBDT) Mode:** FNBDT mode is an enhanced secure voice mode for narrowband connections (RF/Wireless). This mode will allow for interoperability between the STE and Condor (an emerging wireless secure product) secure product lines. Interworking functions required to support the FNBDT Mode have yet to be defined.

(4) **Allied/Coalition Mode:** Yet to be defined.

4.12.3.2 Secure Telephone Unit – 3rd Generation (STU-III)

STU-III equipment is the currently fielded secure telephony device in the U.S. Government. It utilizes LPC-10, CELP, and MRELP (Motorola) algorithms in the 2.4/4.8 and 9.6 kbps modes of operation, respectively. STU-III secure is a voice service implemented on standard PSTN/DSN voice channels. Compression techniques employed on the voice channels must be engineered to support operating algorithms.

4.12.3.3 Advanced Narrowband Digital Voice Terminal (ANDVT)

ANDVT equipment (KYV-5) is primarily used by deployed forces in support of tactical applications. Interface to the GSVS is accomplished through either Radio Wireline Interfaces installed at key communications facilities ashore (NCTAMS) or specifically configured Defense Red Switch Network (DRSN) facilities. ANDVT operations are supported on High Frequency (HF), Ultra-High Frequency (UHF) Satcom, Super-High Frequency Satcom (SHF), and Extremely High-Frequency Satcom. ANDVT utilizes the LPC-10e algorithm. Infrastructure requirements from RF facilities to either RWI or DRSN gateways require 2.4 kbps dedicated digital connectivity per communications channel.

4.12.3.4 Digital Secure Voice Terminal (KY-68)

DSVT (KY-68) is primarily used to provide secure telephony between commanders in the field and at sea. The Ground Mobile Force (GMF) voice architecture can support both 16 kbps and 32 kbps modes of operation. Bandwidth-constrained Naval Units typically operate at the 16 kbps mode. Primary interface to

the field is accomplished by relay via Standard Tactical Entry Points (STEP). Direct interface from the ship to the field is possible if direct satellite and/or Digital Wideband Transmission System (DWTS) is available. Shore interface requires 16 kbps NRZ signaling, while end terminal equipment requires Conditioned Di-Phase (CDI) interface.

4.12.4 Assumptions

- National Security Agency STE and Condor initiatives will continue to drive secure telephony architecture.
- Joint Staff Global Secure Voice Network vision will continue to provide overarching architectural guidance to maximize infrastructure efficiency and minimize operating cost.

4.12.5 Service Architecture

Figure 4-24 illustrates the topology of the DON ship-to-shore telephone connections. Three aspects of the figure are noteworthy.

- There are multiple RF circuits used to make such connections.
- The majority of telephone service is provided through specifically engineered gateway facilities (i.e., NCTAMS or STEP site).
- Connectivity to tactical telephone networks (i.e., TRI-TAC, MSE, etc.) is extended via these shore gateway facilities.

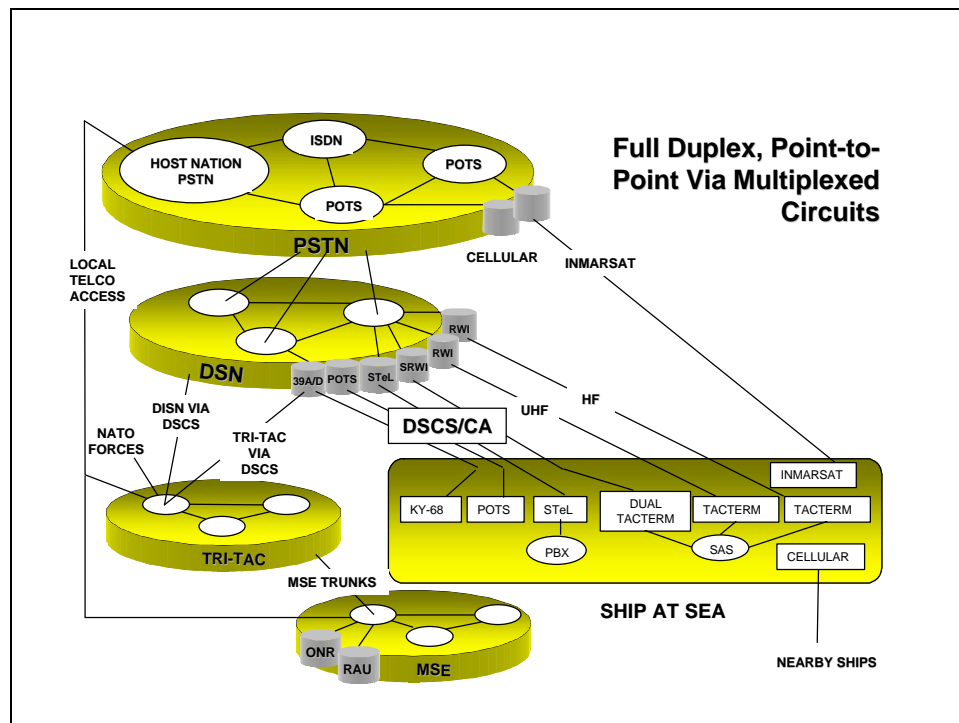


Figure 4-24. Secure Voice Notional Architecture

Figure 4-25 illustrates the components used to provide secure ship-to-shore telephone service. Direct dial secure telephone service is overlaid on basic ship-to-shore telephone connections.

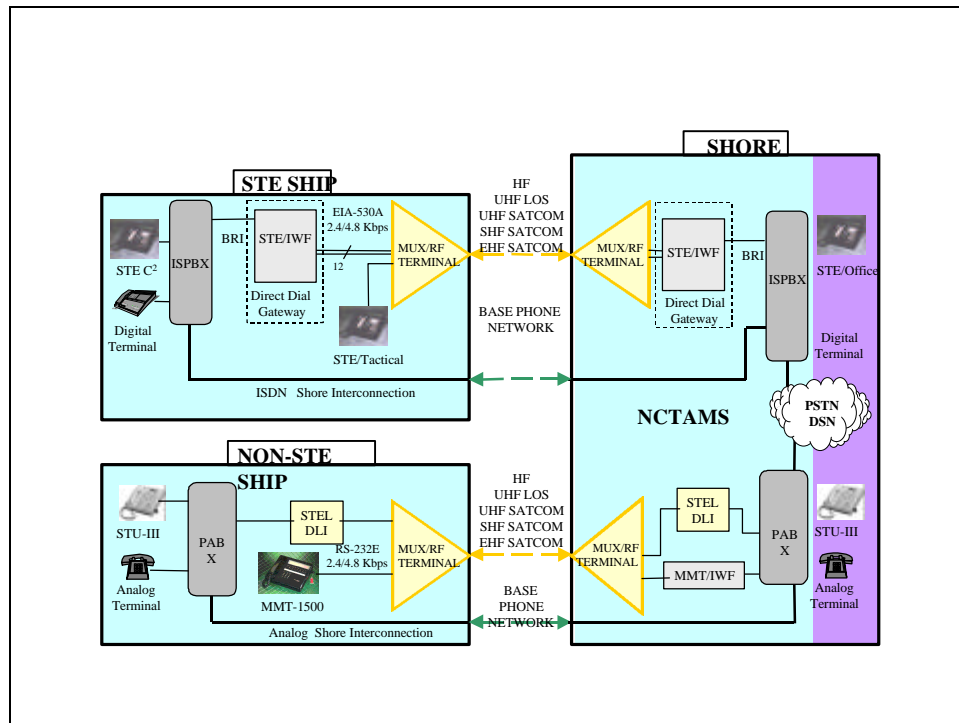


Figure 4-25. Direct Dial Telephony Architecture

Regional Issues and Considerations

Regional planners and designers should ensure that NI-1-compliant ISDN telephone service is available in support of STE secure voice mode.

Campus and Operational Node Issues and Considerations

STE fielding should be a prime consideration in sizing ISDN switches. Initial fielding estimates are being collected by SPAWAR PMW-161.

Deployed Force Issues and Considerations

Existing Indirect DSN access is accomplished via FCC-100 FXS/FXO and Timeplex Voice Server Module (VSM) and FXS/FXO modules. Continued fielding of improved voice modules specifically designed to support STU-III algorithms (I.E. VSM.5) is required to support STU-III until STE Direct Dial is fielded.

4.13 Multimedia

Federal Standard 1037-C defines multimedia as “the processing and integrated presentation of information in more than one form, e.g., video, voice, music, and data.” In this services section,

multimedia services include video teleconferencing (VTC), video applications sharing, video teletraining, and video and image/graphics file servers. Also covered are VTC application-enhanced data services that allow users to share applications and documents and to participate in collaborative activities including video applications sharing, video document sharing, and “white boarding.” Figure 4-26 highlights the technologies supporting the various multimedia services.

| | Video Conferencing | Video Application Sharing | Tele-training | Video/graphics file server |
|-----------------|--------------------|---------------------------|---------------|----------------------------|
| Analog signal | | | | |
| Digital signal | X | X | X | X |
| Real-time | X | X | X | |
| Stored image | | | X | X |
| Point-to-point | X | X | X | X |
| Multipoint | X | X | X | |
| Interactive | X | X | X | X |
| Non-Interactive | X | | X | X |

Figure 4-26. Multimedia Services and Supporting Technologies

4.13.1 Service Description

Video Teleconferencing: Video teleconferencing for Naval forces ashore and afloat allows geographically-dispersed personnel and activities to conduct face-to-face meetings in real time through the transmission of images and sound. Current video teleconferencing systems range from reservation-based, dedicated boardroom systems to portable cart and desktop systems. Desktop video systems based on ATM are emerging. Two transmission models, typical of room-sized and desktop video conferencing, are provided in Figure 4-27.

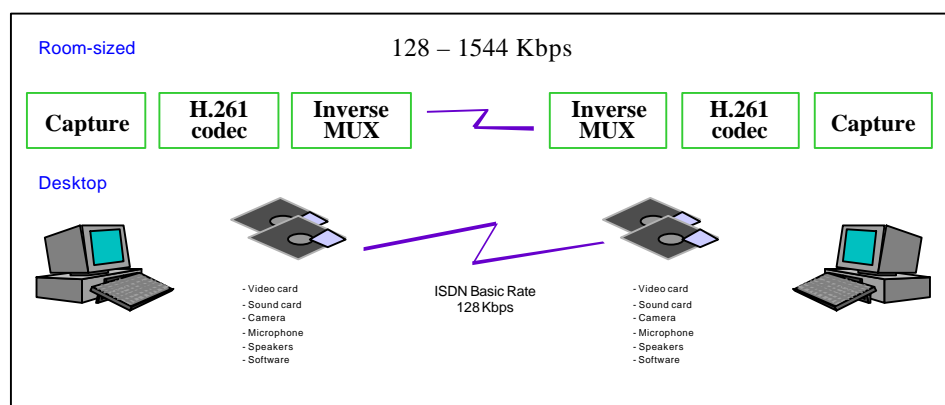


Figure 4-27. Room-sized and Desktop Video Conferencing transmission model

Video Teletraining: These systems are a special class of multimedia services and include interactive two-way video and audio systems, one-way video and two-way interactive audio, one-way video and one-way audio, and multimedia computer-based training applications.

Video File Servers: A video server is essentially a digital storage device (like a large digital video cassette recorder) designed to handle multimedia video content and deliver it to multiple simultaneous users on demand. The video server stores digital video, audio, and graphics in a compressed format that can be retrieved, sorted, and distributed over a communications network.

Multimedia File Servers: Multimedia servers are the more general systems designed to store and forward text, graphics, and images as well as some video and audio files.

4.13.2 Applicable Standards and References

Video Teleconferencing: The video teleconferencing standards in use today were developed by public carriers to promote interoperability between desktop conferencing systems across public network transport service provided by Narrowband Integrated Services Digital Network (N-ISDN). N-ISDN does not specify the network, but does specify the interface to the network. Most existing standards are based on the concept of end user equipment connected to a public network that provides services and connectivity between users.

The DON Information Technology Standards Guidance (ITSG) addresses these end user interfaces and provides an overview of the ITU video teleconferencing standards in ITSG Table 9.5.

Video Teletraining: To be provided at a future date.

Video File Servers: To be provided at a future date.

Multimedia File Servers: To be provided at a future date.

4.13.3 Requirements

Video Teleconferencing: Video teleconferencing must be provided as point-to-point and point-to-multi-point virtual connections for group and personal conferencing. The service must be easy to use and available to be invoked from anywhere. If a desktop video conference session is required, either scheduled or ad hoc, the user or VTC scheduling application simply “calls” another user (or users) to establish a session. The solution must provide for automatic session setup and tear down. The service must use the same transmission medium as voice and data. The service must be secure. The system latency must be low (less than 100 milliseconds) and jitter must be minimal (less than 150 milliseconds) as based upon quality of service guarantees.

Video Teletraining: Video teletraining must provide the following: group or personal interactive two-way video and audio systems, one-way video and two way interactive audio/data, one-way video and one-way audio (broadcast), and multimedia computer-based training.

Video File Servers: Video servers must provide services such as Video On Demand services where the end user has some control over the selection of the material to be displayed and the viewing time. Viewers must have control of the video stream similar to those found on a typical video cassette recorder for restart, rewind, pause, and fast forward. Most Video On Demand systems are point-to-point systems that require some signaling protocol between the end user and the video server. Video servers must also provide encoders to support real-time streaming video; e.g., broadcast video.

Multimedia Servers: Multimedia servers must store and forward text, graphics, and images, video clips, and audio files.

Security Considerations

For persons or activities requiring NSA-approved type 1 link encryption, Figure 4-28 lists the required device(s) for video teleconferencing.

| | ISDN H.320 | POTS H.324 | IP/LANs H.323 | ATM H.321 | Hi-Res ATM H.310 |
|--|---------------|---------------|--|-----------------------------------|-----------------------------------|
| External NSA Approved type 1 Encryption | KG-194 | TBD | KIV-7 at WAN Gateway from Classified LAN | Fastlane Between ATM Nodes; | Fastlane Between ATM Nodes; |
| | KIV-7 | | | KIV-7 at WAN Gateway | KIV-7 at WAN Gateway |

Figure 4-28. Required Link Encryption Devices for Video Teleconferencing

4.13.4 Assumptions

- Naval Service voice, video, and data networks are currently separate, but will share the same transmission medium.
- Video will become a basic network service funded centrally for both capital investment and operations and maintenance.

4.13.5 Service Architecture

There are currently as many methods that provide efficient transport of multimedia conferencing information as there are that provide efficient transport of data traffic. Because few transport methods are capable of meeting the requirements of both, Naval organizations have built separate networks for voice, video, and data. Uncompressed video and audio information requires a high speed Constant Bit Rate (CBR) channel to avoid distorted speech and jerky motion. Data transfers tend to be intermittent or “bursty” in nature with periods of low activity followed by periods of high activity. Figure 4-29 depicts the target multimedia architecture for the Naval services.



Enterprise Issues and Considerations

Interoperability must be ensured through the use of the multi-media/video teleconferencing standards consistent with those used for the Defense Video Service – Global (DVSG). The DON Information Technology Standards Guidance (ITSG) provides detailed information on the DVSG standards.

Regional Issues and Considerations

- Video conferencing services (ad hoc and scheduled) among intra-regional entities or between intra-regional and inter-regional entities are best funded and managed by one using multi-point interpreters/translators.
- At first, all regional video services hubs must accommodate intra-region carrier diversity (e.g., ATM, Frame Relay, Switched Digital) but must transition to a single compatible service over a specified period of time. The target services are based on Asynchronous Transmission Mode (ATM) technology over Category 5 twisted pair or fiber optic cable.
- Multicasting, transmitting IP datagrams to intended recipients, is required for one-to-many or many-to-many applications such as video conferencing, applications sharing, and video teletraining. Existing Ethernet Network Interface Cards (NICs) may need to be replaced by NICs that filter multicasts. This approach prevents forwarding of multicasts to higher network layers (in the protocol stack) for filtering, which will thereby save CPU processing time. In addition, most switches and routers in use today forward all multicast transmissions to all ports, which places unnecessary demands on the network. Network components may need to be upgraded to support several multicast communications protocols that may not be supported in the current configuration. These include:
 - ◆ Internet Group Management Protocol (IGMP)
 - ◆ Distance Vector Multicast Routing Protocol (DVMRP)
 - ◆ Multicast Open Shortest Path First (MOSPF)

- ♦ Protocol-Independent Multicast (PIM) protocol

(While some combination of DVMRP, MOSPF, and PIM are being used today, the long-term goal is to migrate to native multicast using PIM.)

To support time-sensitive audio/video signals, the following additional protocols may be required:

- ♦ Resource Reservation Protocol (RSVP)
- ♦ Real-time Transport Protocol (RTP)
- ♦ Real-time Transport Control Protocol (RTCP)

Because multimedia services use large amounts of bandwidth and typically use the User Data Protocol (UDP) packets that are easier to spoof than standard TCP-based packets, multimedia services should be placed behind firewalls

Deployed Issues and Considerations

Considering the bandwidth demands of multimedia services, stand-alone video teletraining, video file servers, and multimedia file servers is the recommended solution for a deployed shipboard environment.

Remote Shore-based Issues and Considerations

For multi-point video teleconferencing, dial-in via ISDN to a Naval Service multimedia hub is recommended.

For point-to-point video teleconferencing, a dial-in connection can be established if the systems at each end are H.320-compatible or if they use H.323 through a H.320 gateway.

Campus and Operational Nodes Issues and Considerations

- For scheduled board room, rollaway cart, and ad hoc desktop VTC, every system must be H.320-compatible or pass through a H.320 gateway if it uses H.323.
- Because multimedia services use large amounts of bandwidth and typically use the User Data Protocol (UDP) packets that are easier to spoof than standard TCP-based packets, multimedia services should be placed behind firewalls.
- For all new group video-conferencing systems, every system must support the H.261 and H.263 schemas for circuit switched systems to provide potential to use a single 2B+D line instead of three ISDN channels for many applications.

4.13.6 Roles and Responsibilities

Enterprise

- DON multimedia policy, standards, and guidelines are developed, coordinated, and published by the DON Chief Information Officer (CIO).
- The DON CIO reviews waiver requests for the DON CIO multimedia policy, standards, and guidelines.
- Interoperability must be ensured through the use of the multi-media/video teleconferencing standards consistent with those used for the Defense Video Service – Global (DVSG).

- The SPAWAR System Center, Charleston is the DON multimedia/video teleconferencing lab.
- All Naval multimedia/video teleconferencing systems must be certified by the DON.

Regional

- Multi-point Control Units (MCUs) must be placed as close as possible to the enterprise network backbone.
- For multi-point conferences, a Multi-point Control Unit (MCU) must be located somewhere within the network. Figure 4-30 depicts multi-point functionality through the use of multi-point control units.
- Multimedia service providers (Regional/Area/Campus) must establish Service Level Agreements (SLAs) with customers.

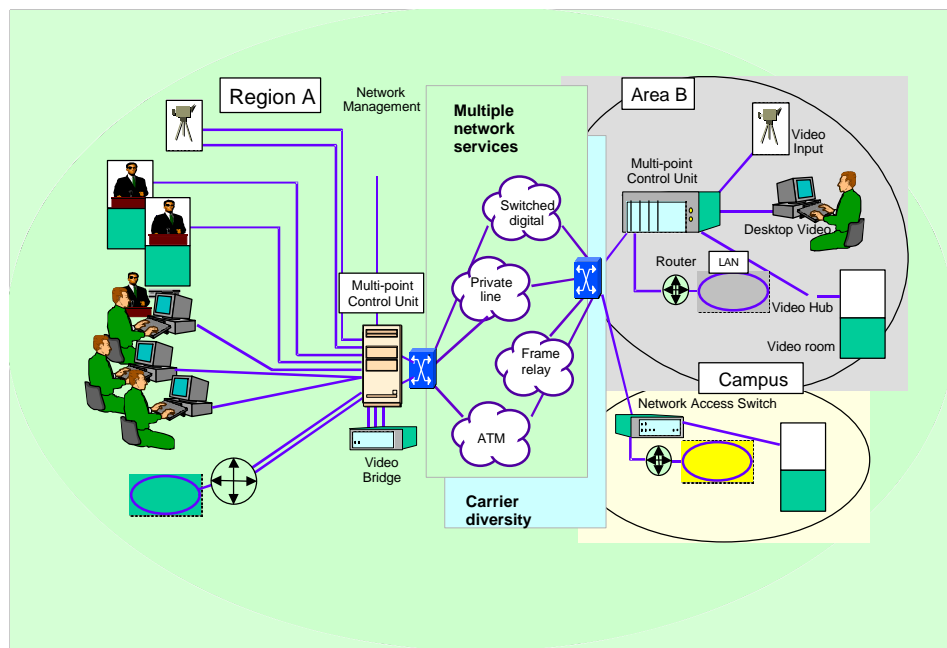


Figure 4-30. Multi-point Control Unit Functionality

Shore-based Campus and Operational Nodes

The ITSC will provide operations support to at least one news server.

4.14 Common Operating Environment Applications

4.14.1 Service Description

The Defense Information Infrastructure (DII) Common Operating Environment (COE) is a set of DoD-wide guidelines, standards, and specifications for the development of software. The DII COE principally provides for software reuse, standardized “look and feel”, and improved interoperability within Joint and Service software-based systems. The DII COE is compliant with DoD *Joint Technical Architecture* (JTA). The DII COE provides a modular open architecture for software development. It has been applied

to software development in the functional areas of C4I, logistics, transportation, base support, health affairs, and finance.

The Office of the Secretary of Defense has issued a directive that all new C4I systems must be compliant with the JTA. The JTA, in turn, mandates the use of DII COE. Combat systems and weapons systems software development will be addressed within future versions of the JTA.

4.14.2 Applicable Standards, Policy, and Guidance

See Joint Technical Architecture (Section 2.2)

See DII COE Integration & Runtime Specifications

See User Interface Specifications for the DII

DISA DII COE references are available at <http://spider.osfl.disa.mil/dii/>

4.14.3 Requirements

Applications developed within the Naval enterprise must abide by the following COE requirements:

- Software development shall abide by the requirements set forth in the DII COE standards.
- Applications developed for operational use within the Naval enterprise must meet, at a minimum, DII COE Level 7 compliance requirements (DII COE levels are discussed later in this section).
- Applications developed for Advanced Technology Demonstrations (ATD), Advanced Concept Technology Demonstrations (ACTD), or Fleet Battle Experiments (FBE) must meet, at a minimum, DII COE Level 5 compliance requirements.

4.14.4 Assumptions

DII COE will continue to be the guidance standards for DoD software development.

4.14.5 Service Architecture

DII COE is structured to provide common services and promote shared database design within DoD-developed software as illustrated by Figure 4-31.

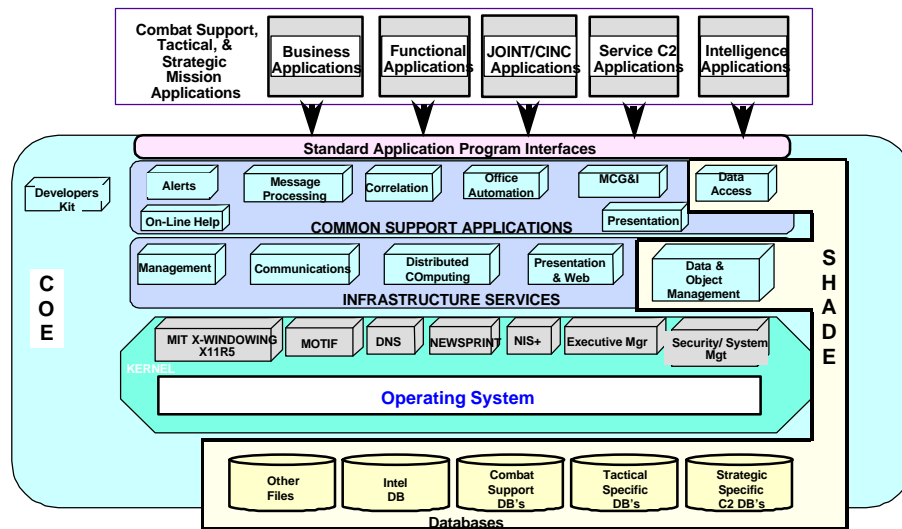


Figure 4-31. General Overview of DII COE Architecture

Adherence to the COE is designated by the depicted eight different levels of compliance. Each higher level of compliance includes all the requirements from the preceding level. Level 8 (Full Compliance) is the optimum level for software applications deployed within the Naval ITI. A minimum of Level 5 (Prototype Compliance) is acceptable for demonstration software which is deployed on a limited basis within the Naval Enterprise such as within an Advanced Concept Technology Demonstration (ACTD), Advanced Technology Demonstration (ATD), or Fleet Battle Experiment (FBE). A minimum of Level 7 (Interoperable Compliance) is the lowest level of compliance acceptable for software integration within the Naval ITI that serves operational (combat and combat support) units.

| DII COE Level | COE Requirements | Naval ITI Requirement |
|--|--|--|
| Level 1 Standards Compliance | <ul style="list-style-type: none"> - Software is based on either NT or POSIX 1.0 compliant OS - Supports DCE, SLIP, PPP, TCP/IP, UDP - GUI is either Motif (with X-Windows) or NT - Supports SQL | Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure. |
| Level 2 Network Compliance | <ul style="list-style-type: none"> - Is COE Level 1 compliant - Supports sockets - Supports DNS, NFS, NIS - Works with C2, BSM security modules - If NT, supports NTFS - Data base transactions implement strict two phase locking | Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure. |

Draft Working Papers of the ITI IPT

| | | |
|---|---|--|
| | <ul style="list-style-type: none"> - Requires no reserved IP addresses - Is not dependent on specific type of LAN | |
| Level 3 Workstation Compliance | <ul style="list-style-type: none"> - Is COE Level 2 compliant - Extensions to OS are clearly documented - Software does not modify associated COTS, GUI, or DBMS - Works with anonymous FTP - Limited "hard-coding" of ports - Software does not modify other software file structures | Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure. |
| Level 4 Bootstrap Compliance | <ul style="list-style-type: none"> - Is COE Level 3 compliant - Software is properly segmented - Software extensions do not conflict with other segments - Segment restrictions are clearly documented - Follows DII COE Directory structure - Software may be installed/de-installed without conflicting with other software segments | Software satisfying this level of compliance is not suitable for deployment within operational portions of the ITI. Deployment is only acceptable within R&D (Software Support Agency) development infrastructure. |
| Level 5 Minimal COE Compliance | <ul style="list-style-type: none"> - Is COE Level 4 compliant - Does not violate COE's UNIX root and login restrictions - Fully compliant with COE Style Guide - Only user interface is via the GUI - No developmental software tools are required - Minimal dependence/conflicts with other segments is present - Inter-segment communication is done via COE structures/processes - Software licensing requirements are satisfied | Minimal level of COE compliance for prototype software deployed within the ITI for ACTD, ATD, and FBE. Deployment is not acceptable within operational portions of the ITI. |
| Level 6 Intermediate Compliance | <ul style="list-style-type: none"> - Is COE Level 5 compliant - Data base access is via COE roles and groups - Segment uses COE web-server - Segment duplicates < 50% of existing COE functions - Segment does not modify environment variables | Segments meeting this level of COE compliance are acceptable for prototype software deployed within the ITI for ACTD, ATD, and FBE. Deployment is not acceptable within operational portions of the ITI. |
| Level 7 Interoperable Compliance | <ul style="list-style-type: none"> - Is COE Level 6 compliant - NT segments use NT registry - Allows cut, copy, paste between segments - Does not replicate data present in Shared Data Environment (SHADE) - Eliminates file permission vulnerabilities | This is the minimum level of COE compliance suitable for software deployment in operational portions of the ITI. |

| | | |
|-----------------------------------|---|--|
| | <ul style="list-style-type: none"> - Rules-based segments have rules within a rules data base - Supports frame-based web services - Does not duplicate COE functions - Less than 25% of COE functions are accessed via private API | |
| Level 8 Full Compliance | <ul style="list-style-type: none"> - Is COE Level 7 compliant - Fully compliant with COE style guide - Uses Joint data elements - No private APIs are used - Uses COE DBMS - Does not duplicate functionality of other segments | This is the objective level of COE compliance for all software applications deployed within the ITI. |

Figure 4-32 COE Levels of Compliance

4.14.6 Roles and Responsibilities

Naval Systems commands and Program Executive Offices are responsible for ensuring that software developed under their cognizance meets the COE requirements in Figure 4-32.

Other Naval organizations which procure or develop non-Program of Record software for inclusion in the Naval Enterprise must abide by the COE requirements in Figure 4-32 whenever possible.

The ITSC planners and engineers should consider the DII COE compliance of software applications within their cognizance to ensure software interoperability within the unit, campus, regional, and global levels.

Volume I, Appendix A – Table of Contents

| | |
|--|------------|
| A. Wide Area Connectivity Plan..... | A-1 |
| A.1 Introduction..... | A-1 |
| A.2 Operating Environment..... | A-2 |
| A.3 Definition of Connectivity and Services | A-2 |
| A.3.1 Wide Area Connectivity..... | A-2 |
| A.3.2 Wide Area Services | A-3 |
| A.4 Acquisition Strategy | A-4 |
| A.5 Functional and Performance Specifications | A-5 |
| A.6 Potential WAN Service Providers | A-7 |
| A.7 Potential WAN Service Providers | A-8 |
| A.8 Evaluating WAN Service Providers | A-10 |
| A.9 Plan of Action for Selecting WAN Service Provider..... | A-11 |

This page intentionally left blank.

A. Wide Area Connectivity Plan

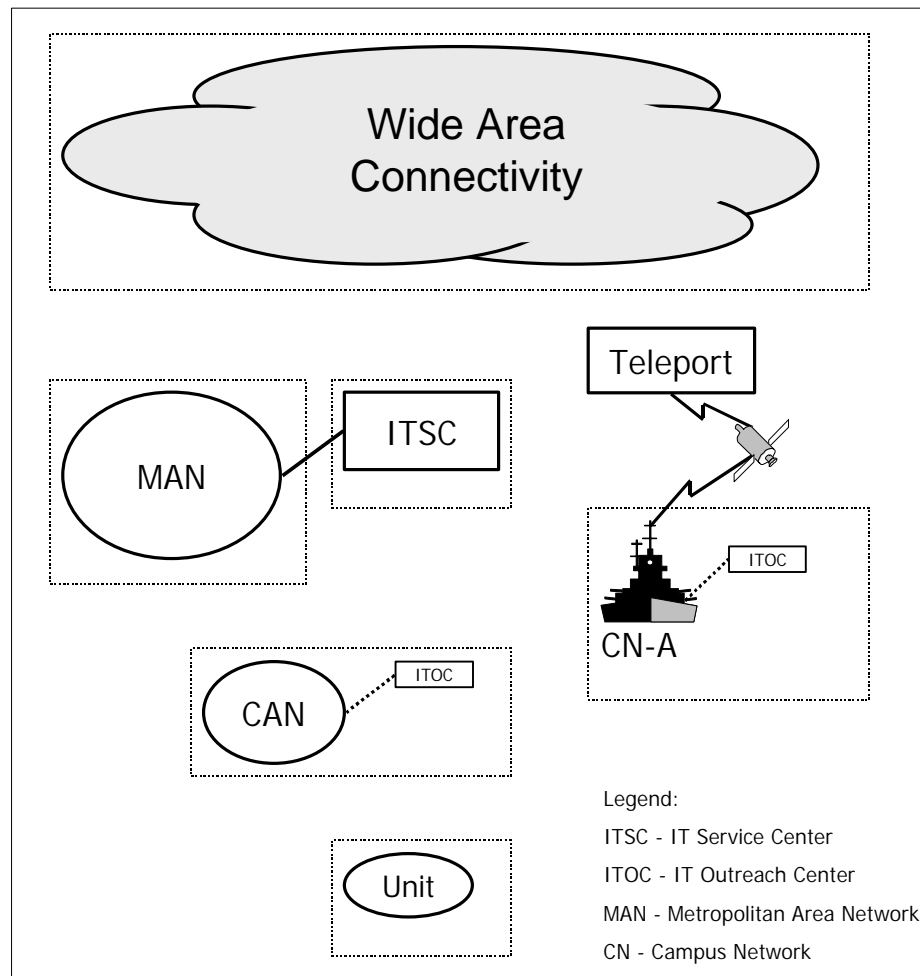


Figure A-1. High Level Components of the TI Architecture

A.1 Introduction

Purpose of the Wide Area Connectivity (WAC) Plan. The principal requirement for the WAC Plan is to outline the strategy and steps to plan and establish a wide area network (WAN) that interconnects all of the DON Metropolitan Area Networks (MANs). Figure A-1 shows the relationship of the Wide Area Network (WAN) to the other components of the Information Technology Infrastructure (ITI) Architecture.

The detail, considerations, and steps in this WAC Plan are consistent with the importance placed on this function by the Navy and Marine Corps. The WAC represents the essential interconnectivity between the organizational elements of the DON enterprise network. Its uniquely complex and inter-related planning requirements necessitate a carefully orchestrated, integrated DON approach.

This WAC Plan provides scope and definition, approach, strategy, and implementation milestones to be used by TI planners to design and implement a DON WAN. This plan will be consistent,

complementary, and interoperable with the overall DON enterprise network as defined by the accompanying planning templates contained in this enterprise network architecture.

This WAC approach is supported by industry best practices and is essential for Naval conformance with the Clinger-Cohen Act of 1996 and the Office of Management and Budget memorandum 97-16.

A.2 Operating Environment

The WAC Plan must address the connectivity of all Navy and Marine Corps organizations, including joint services and supporting contractors, in the following environments:

- **MAN and Outlying Campuses.** As discussed in the network architecture Section 3.1, the MANs provide efficient and effective connectivity to the Naval concentration areas. The purpose of the WAC is to interconnect these MANs
- **CONUS and OCONUS.** Naval concentration areas include both continental U.S. (CONUS) and outside continental U.S. (OCONUS) areas. OCONUS examples include Naples, Bahrain, and Yokosuka. This WAN document initially places emphasis on but does not limit itself to the CONUS areas.
- **Ashore and Afloat.** The terrestrial WANs provide the trunk technology and switching to interconnect the Metropolitan Area Networks (MANs), including afloat units at piers.

For underway and deployed forces, the architecture guidance for the global connectivity provided through satellite communications (SATCOM) and Line of Sight (LOS) radio links is provided separately.

A.3 Definition of Connectivity and Services

A.3.1 Wide Area Connectivity

Definition. WAC provides connectivity that interconnects widely-separated components of the Naval enterprise by interconnecting the Naval MANs. The WAN complements the other network connectivity components and provides end-to-end connectivity across the Navy and Marine Corps units.

Points of Connectivity. The WAN includes long-haul circuits and the ATM switches. The WAN provides connectivity between the MANs (and in some instances, directly to some outlying campuses/bases). As shown in Figure A-2, the WAN includes delivery through shore-based, satellite, at-sea, and other mobile platforms. It addresses connection to the external networks – Defense Information System Network (DISN), Non-secure IP Router Network (NIPRnet), Secure IP Router Network (SIPRnet), Defense Research Engineering Network (DREN), Defense Switched Network (DSN), and the Internet.

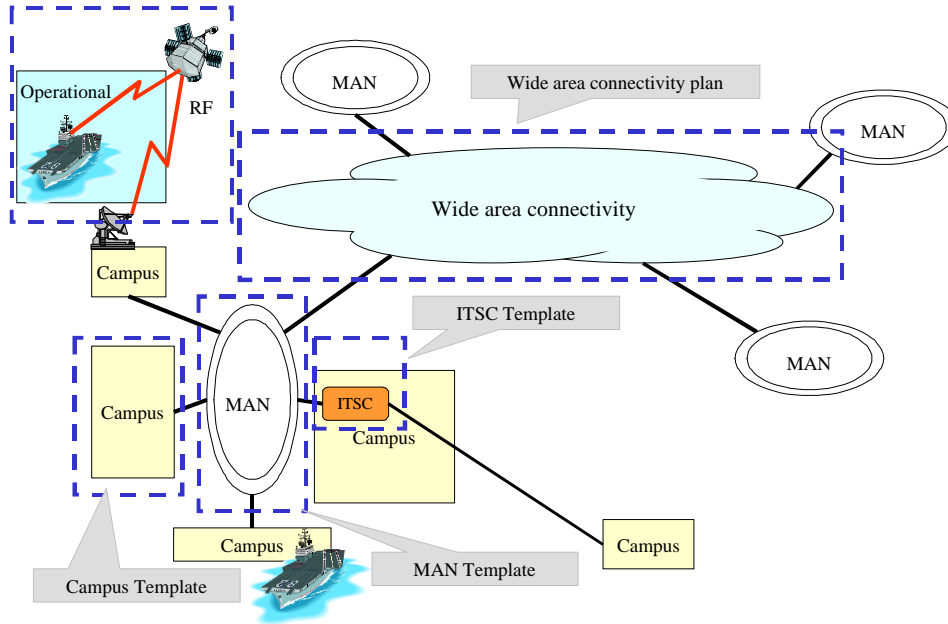


Figure A-2. Wide Area Connectivity Components and Demarcations

Demarcations. The major demarcations for the WAN are the locations of the Metropolitan Area Networks (MANs) and the Information Technology Service Centers (ITSCs). These MANs and ITSCs represent the geographic locations and levels of the DON organization at which most technology infrastructure services planning, implementation, and operation are performed. The MANs and ITSCs figure strongly in determining the DON WAN service functionality and performance.

In specific instances, the ITSCs are the fleet teleports that directly link the WAN to the fleet using a variety of interworking components including satellite and line-of-sight RF links to extend the enterprise network to overseas, deployed, and embarked units.

A.3.2 Wide Area Services

The WAN infrastructure relies upon Asynchronous Transfer Mode (ATM) as its principal technology. ATM enables the consolidation of voice, video, and data onto a single network and offers the advantages of speed, economy, and interoperability for the entire spectrum of networking applications.

Physical Layer – The WAN must support the interconnection of the ATM services in the MAN and campus networks in a way that preserves an end-to-end ATM cell delivery and signaling service. It is expected that the physical layer will be provided by Synchronous Optical Network (SONET) transmission systems for the carrier ATM networks that interconnect the Naval MANs.

Network Management – The MANs and subordinate networks will be managed by the ITSCs. The WAC provider will manage the backbone connectivity, but must also provide enough visibility of the WAC to the ITSCs for DON network managers to eliminate ambiguities and/or isolate failures.

The functions of network management are grouped in three categories: those performed by the WAN service provider, those provided by DON enterprise network management centers, and those performed by the ITSCs.

- **Service Provider Services** - These functions are listed by the name of their top level of granularity. Selected guidance is included for certain functions. The WAN functions provided by the service provider include the following:
 - ♦ Network monitoring - The WAC provider must provide sufficient visibility of SNMP monitoring information to allow DON ITSC managers to correctly isolate faults to the correct provider.
 - ♦ Configuration management
 - ♦ Routing management - The WAC provider must minimize routing hops in the virtual circuits that interconnect the MAN routers and switches.
 - ♦ Trouble response
 - ♦ Performance management - The WAC provider must provide a stated bandwidth and assurance that this can be easily scaled to meet requirements.
 - ♦ Fault management - The WAC Plan must explain how trouble ticket support will be performed. The respective responsibilities of the WAC trouble desk versus the ITSC trouble desk require definition. These should be consistent across the enterprise.
 - ♦ Security management
 - ♦ Network metrics
- **DON Enterprise Management Services** - The functions performed by the DON enterprise management centers are those of an administrative service nature and include ATM and IP address management. The DON Network Information Center (NIC) administrates the ATM addressing plan and the end system addresses. The NIC is also the DON single point of contact for both IP and Network Service Access Point (NSAP) registration. The DON Addressing Plan is provided in Chapter 4. The NIC services are described in the ITSC template (Appendix D).
- **ITSC Services** - The balance of the required services are performed by the ITSCs. These ITSC functions are delineated in Appendix D and Chapter 4.

A.4 Acquisition Strategy

Subscription Based on Functionality. The overall DON strategy for obtaining enterprise backbone services is one of subscription – coordinated, carefully determined, and consistently implemented subscription of services. Multiple DON organizations and multiple commercial service providers may team to acquire the aggregation of WAN services that provide the complete, robust, and efficient interconnectivity required by all Navy and Marine Corps organizations. These services extend to units operating in CONUS and OCONUS regions and to deployable units operating at sea, in remote sites, and pier-side.

In the case of CONUS regions, each region provides connectivity to all other regions either directly or through neighboring regions. The collection of regions, the services they provide, and the connectivity between them constitute the DON enterprise network. It is through implementation agreements and the design factors and architecture guidance in this document that the DON regions successfully share responsibility of a well-designed and managed DON enterprise network.

Service Level Agreement. The subscription of services will be based upon very specific service level agreements that include functional specifications as well as performance specifications. This WAC Plan will help provide guidance to participating DON organizations to adequately define

these specifications. This will ensure that the required functionality, performance, and business case analysis guidance is specifically stipulated in resulting Request for Proposals and in subsequent contracts that define these services.

Competitive Strategy. Determining functionality, performance, and cost (business case analysis) factors of the backbone services requires multi-disciplinary skills for planning and implementation. The responsibilities for implementing the WAC plan include development of the specification and design, execution of the acquisition and funding, and operation and maintenance. The WAC Plan will provide a context for evaluating and selecting the factors that pertain to these three responsibilities.

A.5 Functional and Performance Specifications

In most cases, the WAN is described not in terms of design guidance, as is the case of MAN and campus templates, but in terms of service levels or outcomes. In other words, the government will specify the required services but not the specific implementation architecture. Exceptions to this will occur only when specific technologies are required for functionality or interoperability. For example, the architecture will not specify SONET but will require ATM. The service provider will be allowed to use the architecture and technology solutions that best achieve the specified services and provide the best business practices cost. The specifications are stated in a form that provides an empirical level of expectation and description of the services to be provided.

The services provided by the WAN are in five categories: security, functionality, interoperability, performance, and cost. The level of service, and in some cases, the specific technologies required, are described for each. These specifications are provided as a guide; actual requirements may warrant adjustment of these values but this should be based on solid documentation.

- **Security** - Few security mechanisms are allocated to the wide area connectivity. For the protection of classified information, cryptographic equipment is required. For the security of the infrastructure, the solution is to provide redundancy in the network components. Redundancy should be used when available. The use of Simple Network Management Protocol (SNMP) and remote monitoring of critical networking components such as switches, multiplexers, and routers should be emphasized and control of these is warranted whenever possible. This architecture details the appropriate mechanisms that must be supported.
 - ♦ Hardened
 - Level of service: must provide a sufficient level of protection for denial of service and intrusion detection.
 - Technology: N/A
 - ♦ Survivability (specifically relating to security) (see also Performance)
 - Level of service: intrusion detection, denial of service, and vulnerability to service attacks. The WAC provider is responsible for these only as they relate to the WAC-provided service. They will normally consist of data link layer encryption services (e.g., Fastlane).
 - Technology: N/A

- **Functionality** - The capabilities required of the network infrastructure that are necessary to effectively and efficiently support the operational mission and requirements must be clearly defined.
 - ♦ ATM
 - Level of service: support voice, video, and data.
 - Technology: end-to-end switched virtual circuit (SVC) and permanent virtual path (PVP).
 - ♦ MAN connectivity
 - Level of service: each MAN is provided access to two geographically-separated WAN switches. This is directly related to Availability (Ao) and Survivability.
 - Technology: N/A
 - ♦ Switch functionality
 - Level of service: must support constant bit rate (CBR) QoS.
 - Technology: non-blocking and provide separate queuing for different Quality of Service (QoS) classes. WAN switches will support at least 2048 switched and/or permanent virtual circuits (PVCs) per interface.
 - ♦ Availability
 - Level of service: accessible to all MANs and outlying bases. (See also Network Availability for additional Availability requirements.)
 - Technology: N/A
- **Interoperability** - The components of the network infrastructure must interconnect and efficiently and effectively communicate signaling and other information transfer data.
 - ♦ Service Delivery Points (SDPs)
 - Level of service: must support the full suite of ITSG-cited ATM protocols and addressing schemas. The WAC SDP must be able to provide the required support of DON IP requirements as well as the cited Interior Gateway Protocols, Exterior Gateway Protocols, and firewalls.
 - Technology: N/A
- **Performance** - The WAN service must be sufficiently qualified to meet the following information transfer requirements to support the Naval mission:
 - ♦ Bandwidth
 - Level of service: minimum of OC-3.
 - Technology: redundant dual-homed.
 - ♦ Delay
 - Level of service: maximum end-to-end delay for CBR service connection should be less than 200 us/hop for processing and queuing time plus the necessary propagation delay (5 us/km). The maximum end-to-end cell delay variation should be less than 1 ms.

- Technology: N/A
- ♦ Latency
 - Level of service: if switched point-to-point SVC or SVP is supported, the maximum latency in completing a call setup must be less than 100 ms/hop plus propagation delays.
 - Technology: N/A
- ♦ Network availability
 - Three principles of high availability are established: (1) eliminating single points of failure, (2) reliable crossover, and (3) prompt notification of failures as they occur. All three must be accounted for in the performance metrics.
 - Network monitoring is an issue that requires visibility. A problem in monitoring network availability is determining the failure point when something breaks. The normal situation is for the commercial vendor and base telecommunications to each deny responsibility. Through the ITSC, the capability should exist to immediately sort out the failure point and to call the correct repairman. One means is SNMP visibility of the vendor's network—having access in the ITSC to the real-time availability data that the vendor system monitors are indicating. Having this access should be part of the WAN service provider specification.
 - Level of service: 99.99 percent.
 - Technology: no switch or physical circuit is a single point of failure.
- ♦ Survivability (specifically relating to performance - see also Security)
 - Level of service: vulnerability to forces of nature and human action, including enemy action (i.e. backhoe, power loss, terrorist strike).
 - Technology: no switch or physical circuit is a single point of failure. It is strongly recommended that the service provider's assumptions for survivability be carefully investigated and validated.
- ♦ Quality of Service
 - Level of service: ability to support a number of service classes based on the traffic type, each with an associated QoS parameter.
 - Technology: N/A
- **Cost** - Implementation cost is an important category for judging value and must be done in the context of the level of service.

A.6 Potential WAN Service Providers

The DON CIO policy for subscription of WAN services is based on a competitive strategy. The functionality and performance of the service must support that required by the individual MANs and be provided at a competitive price. When DISA can provide the same or equal service at the same or less cost, they will be the preferred service provider. When DISA is not competitive with alternative service providers, subscription will be obtained elsewhere based on best value.

A.7 Potential WAN Service Providers

Three service providers are suggested for consideration. They each have specific strengths and weaknesses that must be weighed in a WAN decision. The WAC Plan could use the services of one or all three, or there may be other more acceptable approaches such as a DON-administered contract to a commercial provider.

DISA

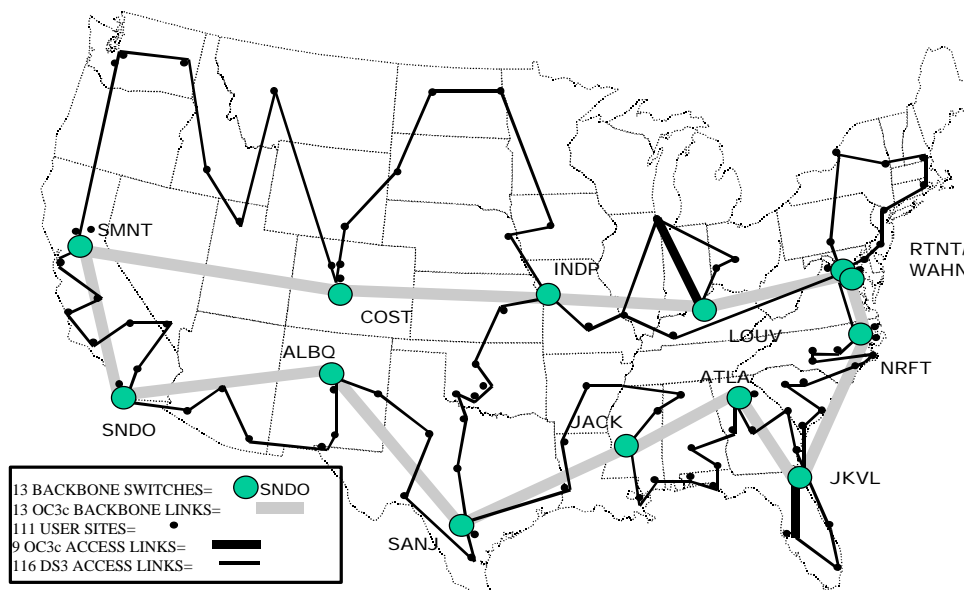


Figure A-3. DISA Wide Area Network Connectivity

The top level view of the Defense Information Systems Network (DISN) in Figure A-3 shows the CONUS ATM service provided by DISA. DISN services in the OCONUS environment represent a near global coverage for Navy and Marine afloat, combat ground, and ashore operating units. In July 1998, 50 ATM nodes were available globally for DoD units. By October 1998, the number of nodes was estimated at 125. The projection for December 2001 is 200.

SmartLink

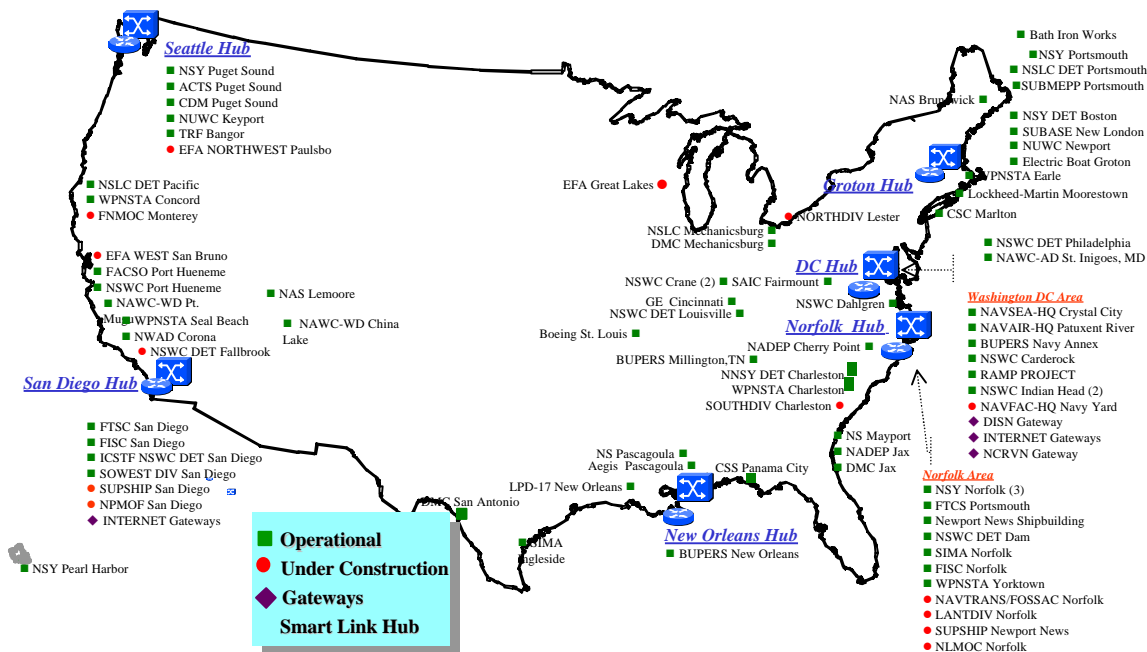


Figure A-4. SmartLink Wide Area Network

The top level view of the SmartLink Backbone Plan depicted in Figure A-4 shows SmartLink's long-haul SONET circuits, ATM switches, and wide area supporting services. This backbone plan is aligned with the fleet concentration areas described in Chapter 2-4. Accordingly, the plan provides an ATM switch at each of the identified concentration areas. SmartLink provides connectivity between the regions (and in some instances, directly to some outlying campuses/bases).

DREN

Defense Research Engineering Network

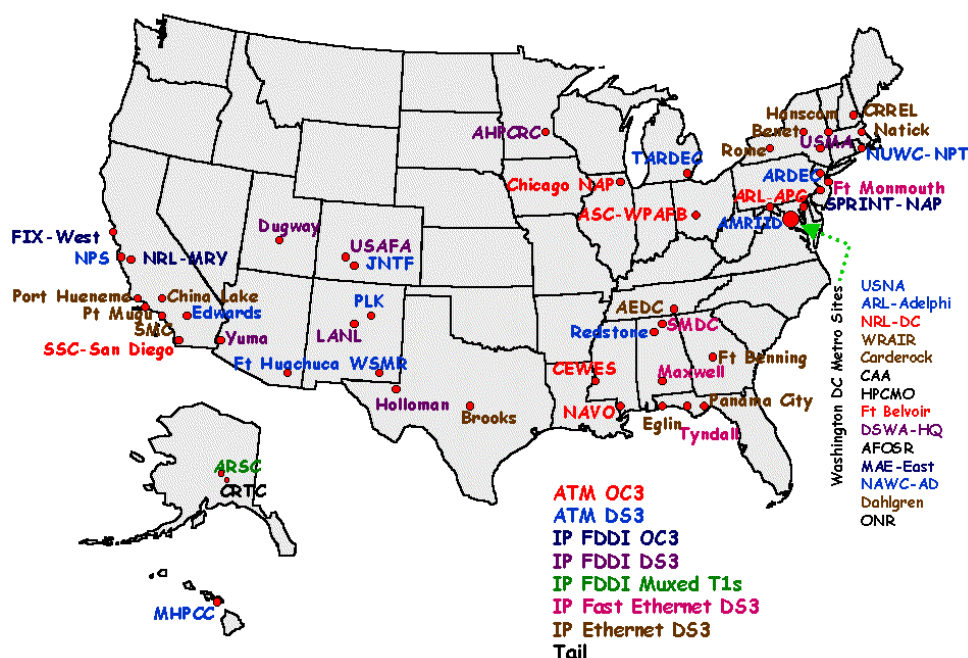


Figure A-5. DREN Wide Area Network

Defense Research Engineering Network (DREN) is a DoD WAN used to link DoD scientists and engineers. The DREN project provides wide-area ATM networking services at bandwidths that are commensurate with DON user communication requirements. ATM connectivity currently includes over 20 CONUS locations.

A.8 Evaluating WAN Service Providers

The WAN services source determination will be made by a DON Integrated Product Team (IPT) comprised of representative (Navy and Marine Corps, tactical and non-tactical, afloat and ashore) operational and functional experts.

The determination will be made based on a balanced scorecard approach using as a basis the characteristics defined in the functional and performance specifications in section A.5. Addressing the proposals of each competing WAN service provider will be performed in accordance with the criteria in Figure A-6.

| Network Characteristic | Raw Score | Weighting (.xx) | Adjusted Score |
|------------------------|-----------|-----------------|----------------|
| Network Security | X | .30 | X |
| Functionality | X | .25 | X |
| Interoperability | X | .25 | X |
| Performance | X | .20 | X |
| Total | — | 1.0 | X |
| Total Merit / Cost | — | — | X/\$YY |

Figure A-6. Service Provider Evaluation

Four of the five characteristics will be scored based on a numerical rating system (1-10). The supporting subsets of each characteristic will be evaluated and aggregated to determine the value of the characteristic.

Also, each characteristic will have a weight factor based on its judged importance to the overall Naval mission. The total of these individual weighting factor will be 1.0. Recommended weighting factors are provided.

The numerical value for each characteristic is multiplied by its weight factor and the total of the five adjusted scores equals the relative merit for the source provider.

Cost will be viewed as a separate variable. The relative merit will be expressed in terms of a ratio of total merit to cost (e.g, X / \$ YY).

A.9 Plan of Action for Selecting WAN Service Provider(s)

The DON CIO Board of Representatives will charter an IPT or task a standing IPT to select a service provider(s) of DON's Wide Area Network services. This team will include individuals with appropriate acquisition experience.

The IPT will evaluate alternative sources to meet the DON enterprise WAN requirements as defined by DON mission requirements. The solution will be consistent with this DON ITI Architecture document and the DON Information Technology Standards Guidance.

The WAN services to be provided must consider wide area connectivity today, what is required, and consider the best way to resolve the service gap. The resulting strategy must identify specific:

- Sites within MANs and those outside

- Sites for which service points of presence exist and those outlying sites that have significant investment, implementation, or provider issues (tail circuits, OCONUS remote sites, and technology change over)

The IPT will use the outline of the balanced scorecard to develop a data sheet for the purpose of going to DISA, SmartLink, and DREN (and potentially others) to ask for a prospectus of WAN services provided via Request for Proposals (RFPs). The RFP shall place emphasis on obtaining complete data so that meaningful comparisons can be made.

The IPT will complete an analysis of the alternative providers and present a recommendation in writing to the Board of Representatives. Included in the presentation will be a recommendation for the administration and funding of the WAN.

Volume I, Appendix B - Table of Contents

| | |
|---|------------|
| B. Metropolitan Area Network Design Template | B-1 |
| B.1 Purpose | B-1 |
| B.2 Drivers | B-2 |
| B.3 Elements, Features, and Specifications | B-3 |
| B.3.1 ATM Overlay | B-3 |
| B.3.2 IP Overlay | B-7 |
| B.3.3 MAN Specifications | B-9 |
| B.4 Service Considerations | B-10 |
| B.4.1 ATM MAN Ownership..... | B-10 |
| B.4.2 Selecting a Commercial Service Provider..... | B-11 |
| B.4.3 Balanced Implementation | B-12 |
| B.5 Outcome Based Implementations - Metrics | B-13 |
| B.5.1 Security | B-13 |
| B.5.2 Functionality | B-14 |
| B.5.3 Interoperability | B-15 |
| B.5.4 Performance | B-15 |
| B.5.5 Cost..... | B-16 |
| B.6 Evaluating MAN Products and Services | B-16 |

This page intentionally left blank.

B. Metropolitan Area Network Design Template

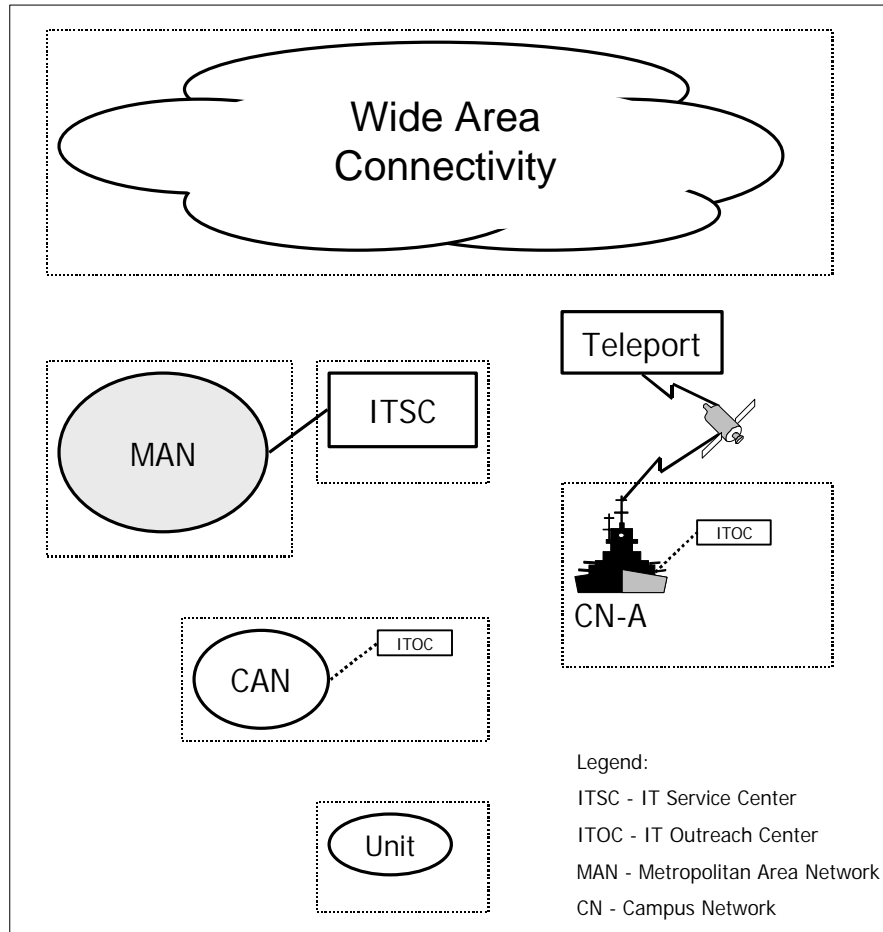


Figure B-1. High Level Components of the TI Architecture

B.1 Purpose

The foundation for integrated, enterprise IT and systems in the DON is the technology infrastructure. Figure B-1 depicts the relationship of the Metropolitan Area Network (MAN) to the other components of this infrastructure. This MAN template defines the technology infrastructure interfaces between the Base, Post, Camp or Station, their associated Information Technology Service Center (ITSC), and any other entities that may constitute a given region. It also addresses connectivity to the DON enterprise network or WAN.

The MAN template is intended to provide design guidance to infrastructure planners and implementers. This guidance is provided in the form of connectivity graphics and associated relevant descriptive text. The template supports decentralized design and implementation of connectivity solutions that are consistent, complementary, and interoperable with the overall DON technology infrastructure.

The MAN approach enables the consolidation of voice, video, and data onto a single network and offers the advantages of speed, economy, and interoperability for the entire spectrum of networking applications. The MAN's underlying physical infrastructure connectivity is essentially leased by the DON regional managers from local and regional wide area service providers. Riding on this leased network is a DON enterprise infrastructure that relies on ATM. ATM is the technology basis for the Navy and Marine Corps autonomous networks and community of interest networks.

B.2 Drivers

The DON mission requirements require a world-class information infrastructure. Each MAN is an integral part of the DON enterprise network and supports the collective requirements of the Navy and Marine Corps customers operating within an individual fleet concentration area. The set of drivers that must be addressed is:

- ***Information superiority*** - This architecture must support improved security, functionality, performance, configuration management, and cost avoidance and provide cheaper, better, and faster service.
- ***Decentralized implementation*** - The DON campus and base networks are implemented in a decentralized fashion and must be based upon a clear and well-defined technical architecture that provides a basis for interoperability in the region, between campuses, and with external organizations.
- ***Customer-based*** - The connectivity and services required by all DON customers operating in a metropolitan area (or region) must be provided for by a fully functional MAN.
- ***Consolidation*** - Underused or common information technology infrastructure (ITI) functions must be consolidated and streamlined whenever appropriate to provide better services at significantly less cost.

B.3 Elements, Features, and Specifications

This section describes the technology elements of the MAN template, their operational features, and their design specifications. The specific sections are depicted in Figure B-2.

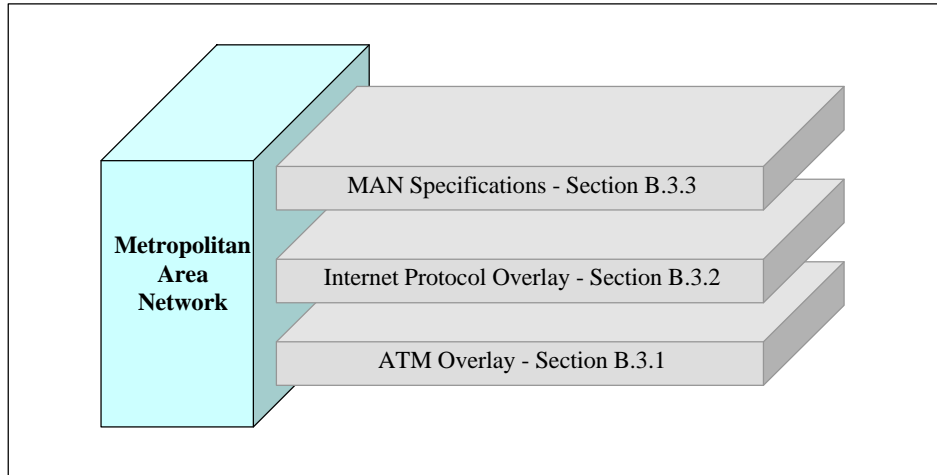


Figure B-2. Layered Description of Metropolitan Area Network

The guidance provided here is intended for technical planners and implementers to produce MANs that are based on the best technology solutions available, that offer satisfactory levels of network performance and reliability, and that can operate successfully in the context of a global enterprise network. This section will require frequent updating to ensure that changes in technology, services offered, and Naval requirements are adequately reflected.

B.3.1 ATM Overlay

The notional MAN template in Figure B-3 shows typical DON campuses or bases that are connected to the MAN backbone using ATM switches. The features of the ATM overlay are described in subsequent subsections titled Physical Layer, Switching and Routing, and Security.

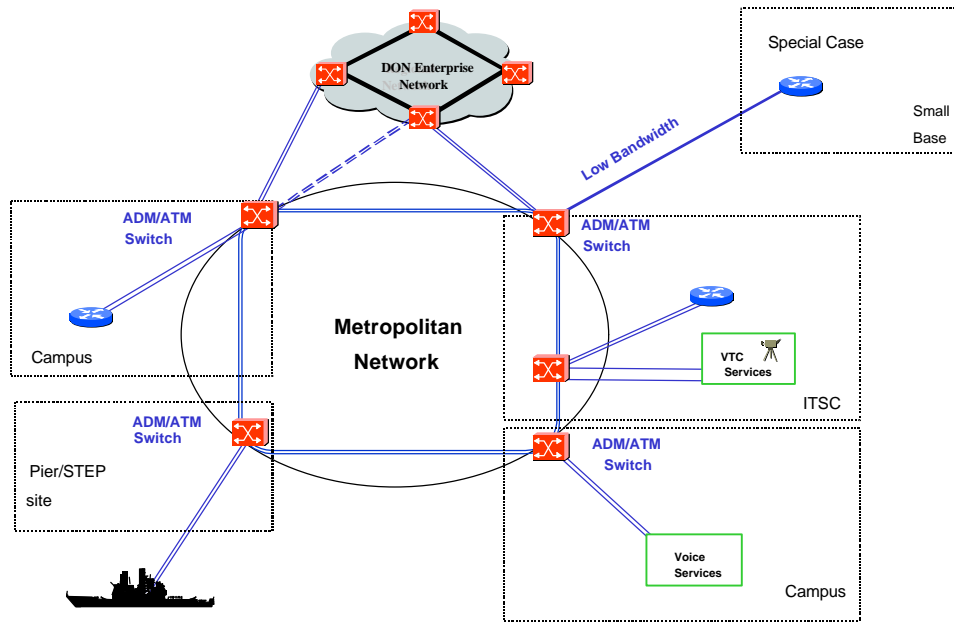


Figure B-3. Metropolitan Area Network Template (ATM Overlay)

B.3.1.1 Physical Layer

The MAN switches are connected using an underlying transport alternative that can be a point-to-point mesh using dark fiber, a point-to-point or ring configuration using Synchronous Optical Network (SONET) technology, or an ATM regional network. For any of these three methods, the conditions under which the service provider supports the underlying DON ATM MAN is important to establish. In any support arrangement, the DON must maintain positive control over the MAN ATM infrastructure to ensure that the transport technology is protocol-independent and to ensure support of a variety of ATM configurations that operate on top of the physical layer.

It is important that the physical layer support high network availability. This requirement necessitates dual threading (redundancy in components and circuits) to thereby eliminate single points of failure.

B.3.1.2 Switching and Routing

Voice, video, imagery, and data are encapsulated in ATM cells and transported over the MAN backbone. The underlying MAN transport is transparent to the users of the campus and base networks as long as the service exhibits satisfactory availability, bandwidth, and latency. A number of configurations can be supported between the MAN and the campus environments including switch-to-switch links and switch-to-router links. The most desirable configuration is the switch-to-switch link (with ATM switching available at each campus endpoint); the switch-to-switch link provides the most flexibility and robustness in satisfying networking requirements.

ATM switch connections to the carrier backbone should be dual-homed for appropriate redundancy as shown in Figure B-3. ATM switch connections may also be the preferred solution for campus sites that use the MAN backbone to interconnect voice switches (e.g., 5ESS) to the DON infrastructure. Aggregation of campus video teleconferencing infrastructures (using multiple Basic Rate Interface (BRI) Integrated Services Digital Network (ISDN) channels) into the MAN ATM

switching fabric is also desirable to the extent that the ATM backbone can be leveraged to provide cost savings over FTS-2000.

B.3.1.3 Security

The placement of security mechanisms in the ATM overlay as shown in Figure B-3 is highly dependent on the type of MAN construction. Two basic cases for security solutions are presented in this section: MANs using commercial ATM services (not under positive DON control) and MANs using point-to-point fiber or SONET links (or mesh) under DON-controlled ATM services.

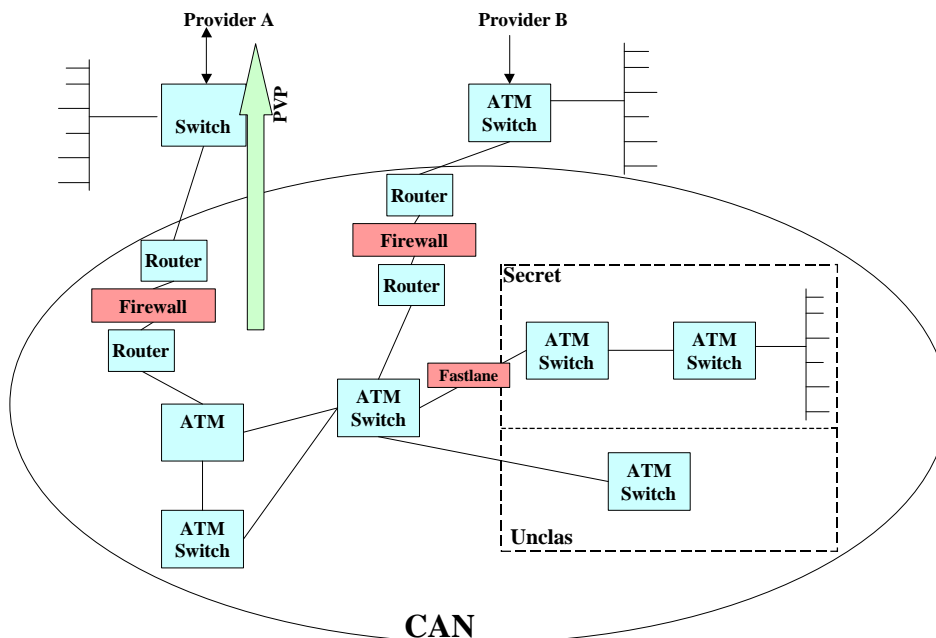


Figure B-4. MAN ATM Not Under DON Control (Security Case 1)

In the case illustrated in Figure B-4, the commercial ATM service provider and other customers receiving services from the commercial ATM cloud must be assumed to be threats to the confidentiality, integrity, and robustness/reliability of the DON ATM overlay. To mitigate these threats, the following security mechanisms are required:

- Fastlane ATM encryption devices must be employed to encrypt the SBU information passed between campuses. Secret information is already encrypted before it enters the CAN (at higher layers in the ISO model using object level security, such as secure e-mail) via Virtual Private Networks or under NES/EIP enclaving.
- The vendor's ATM cloud must be analyzed in detail in order to determine its susceptibility to single points of failure. If the commercial cloud is not found to have sufficient redundancy, a second provider must be used to provide redundancy. This points out the need to have visibility into the vendor's management domain. One way to obtain this visibility is to export on a continuing basis the relevant Simple Network Management Protocol (SNMP) data from the vendor to the ITSC.

- A Permanent Virtual Path (PVP) solution, with appropriate committed information rates, will be used to ensure that the commercial provider can satisfy the minimum bandwidth requirements of the MAN. However, if DISN ATM services are subscribed, it may be possible to use a Switched Virtual Circuit (SVC) solution with Memorandums of Agreement (MOA) guaranteeing a minimum bandwidth.
- Bandwidth allocation management within the DON-controlled portion of the ATM overlay must be provided. This portion of the overlay may actually exist in the CAN when commercial ATM services are used to construct the MAN. This management must provide administrators with the capability to identify the priority of data transport requests and to allocate network bandwidth to the highest priority when contention occurs. Additionally, the ATM overlay must feature mechanisms to “order and add” additional required bandwidth from the commercial ATM cloud.
- The ATM overlay must provide mechanisms that ensure that the DON-controlled components of the overlay (these may be considered part of the CAN) can only be managed by authorized administrators, are resistant to penetration attempts, and are resistant to ATM signaling based denial of service (DoS) attacks. The use of KG-75 Fastlanes significantly reduces the potential for unauthorized administration and successful penetration originating from outside the DON-controlled portion of the overlay. In order to reduce the potential for successful penetrations originating from within the DON-controlled portion of the overlay, remotely-managed network components must feature a non-spoofable authentication mechanism.

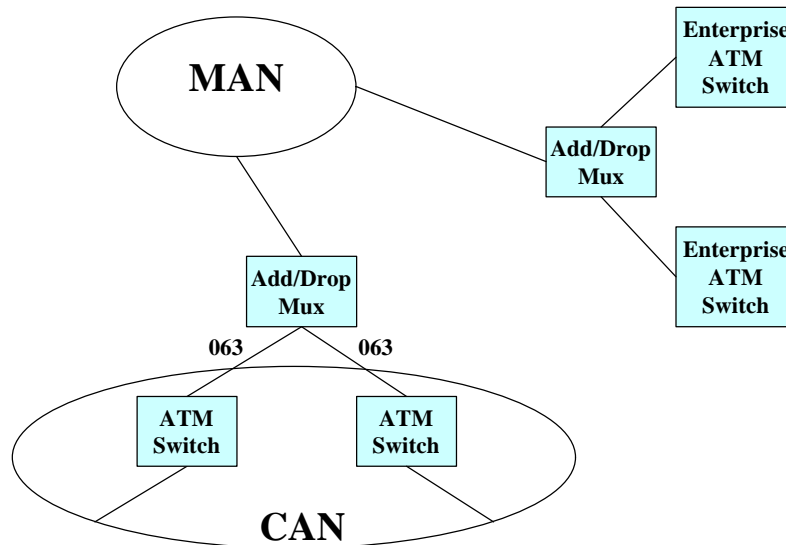


Figure B-5. SONET MAN Under DON Control (Security Case 2)

Figure B-5 shows the condition of implementing a SONET MAN service in which the MAN is under DON control. The potential threats to the confidentiality, integrity, and survivability of the DON ATM overlay are significantly reduced. However, certain security mechanisms are required and are outlined in the following.

- Fastlane ATM encryption devices are only required for secret information (installed between secret buildings and SBU CANs) because the entire ATM MAN is operated at the SBU system high level.
- The ATM MAN must be constructed in a redundant fashion such that the failure of a single network component or interconnection will not lead to the disconnection of a CAN from the MAN or failure of the MAN.
- Bandwidth allocation management must be provided within the ATM overlay. This management must provide authorized administrators with the capability to identify the priority of data transport requests and allocate network bandwidth to the highest priority when contention occurs. Additionally, the ATM overlay should be designed with the ability to “order and add” additional bandwidth as required (such as by adding additional links to the ATM mesh).
- The ATM overlay must provide mechanisms to ensure that DON-controlled components of the overlay (these may be considered part of the CAN) can only be managed by authorized administrators, are resistant to penetration attempts, and are resistant to ATM signaling-based DoS attacks. In order to reduce the potential of successful penetrations originating from inside the DON-controlled portion of the overlay, remotely managed network components must feature a non-spoofable authentication mechanism.

B.3.2 IP Overlay

The MAN IP overlay describes the connectivity infrastructure which transports IP traffic onto the MAN and then through Navy and Marine Corps organizational sites. This MAN architecture spans the physical, network, and application layers and leverages ATM to support a robust implementation of IP. In combination, the IP and ATM technologies support the creation of a high performance, flexible, and proven infrastructure. The features of the IP overlay are described in these subsections:

- Integration with ATM
- Routing Determination
- Performance Provisioning
- Security

B.3.2.1 Integration with ATM

The MAN template in Figure B-6 builds on Figure B-3 and depicts the close relationship between the IP layer and the ATM layer. The IP networking design is irrevocably tied to the ATM implementation. The specific methods for transporting IP over ATM, such as using request for comments (RFC) 1483 encapsulation techniques, RFC 1577 address mapping, and the internetworking capabilities of LAN Emulation (LANE) and Multiple Protocol Over ATM (MPOA), are discussed in Chapter 2. The ATM technology attributes also provide the capability to transfer IP packets with a “Quality of Service”-like guarantee of network performance.

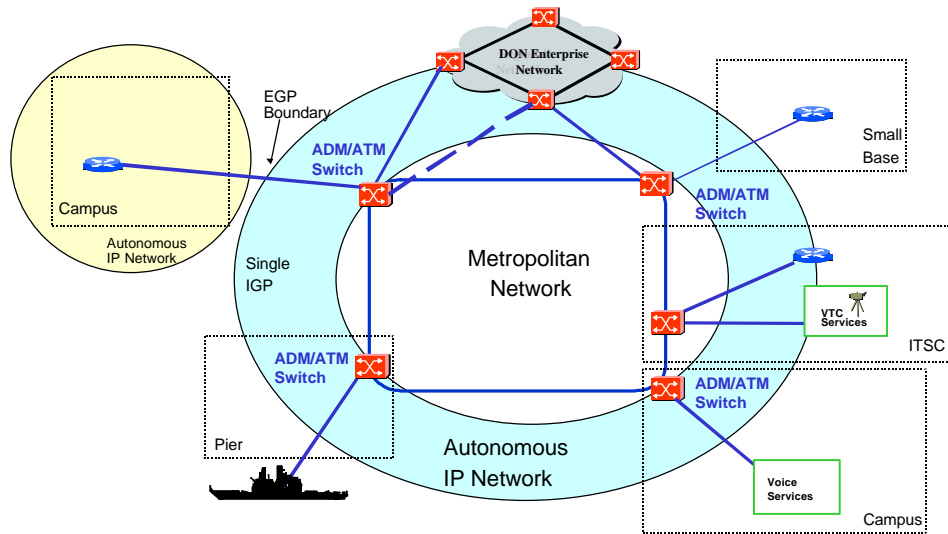


Figure B-6. Metropolitan Area Network Template (IP Overlay)

B.3.2.2 Routing Determination

The general design for internal and external routing is described here as a basis for MAN implementation decision-making. For internal routing, Open Shortest Path First (OSPF) is the preferred Interior Gateway Protocol (IGP) to connect the campus/base to the backbone MAN. There are other satisfactory solutions by which bases can exchange information with the MAN, but only the preferred solution is addressed here.

The MAN will be OSPF Backbone Area “0”, the ITSC is a separate backbone area, and the campus/base can be a separate area.

In the case of the campus/base, the alternative to OSPF is to use an external router solution based on Border Gateway Protocol (BGP) 4 – if supported by a favorable evaluation in performance tradeoffs. The MAN/regional ITSC will provide the OSPF “border router” and will also provide access to other autonomous systems.

The DON autonomous systems as described in Chapter 2 will connect to the MAN using BGP 4 as the preferred Exterior Gateway Protocol (EGP).

B.3.2.3 Performance Provisioning

The MAN template depicts a model using IP and ATM technology that provides for flexible, high-bandwidth, IP-based “autonomous networks” on top of ATM-based “virtual circuits.” This design supports the flexibility to provide bandwidth as it is needed without being constrained by congested IP paths. These services should extend to all campuses within a MAN. In support of the fleet sites, the autonomous network service extends to the piers, to the Naval Computer and Telecommunications Area Master Stations (NCTAMS), and to the Standard Tactical Entry Point (STEP) sites.

B.3.2.4 Security

The DON IP overlay will use the DON ATM overlay to provide required confidentiality, integrity, reliability, and robustness. Confidentiality and integrity for SBU information are provided by ATM on a bulk basis. It is further concluded that for confidentiality, authenticity, and non-repudiation of classified information, additional network and higher layer tools must be provided. The following guidance outlines the additional security mechanisms required to secure the DON IP overlay at the MAN level.

- Connections to external networks (NIPRnet at SBU and SIPRnet at Secret) will only be provided via regional/MAN firewalls. These firewalls will enforce a uniform access control policy for outside NIPRnet and SIPRnet entities attempting access to the DON IP overlay. These firewalls will most likely be installed at regional ITSCs or firewall facilities (FWFs). A centrally-monitored intrusion detection system (IDS) should be installed in concert with the network firewall at all interconnections to external IP networks. Additional information on network firewalls and IDSs are provided in Appendix E.
- Exchange of routing table information updates between DON IP overlay routers will use cryptographic authentication mechanisms as specified in the DON CIO ITSG section 3.4.1.4.
- The IP overlay must provide mechanisms to ensure that network components of the CAN are only managed by authorized administrators. At a minimum, network components that are remotely managed must feature a non-spoofable authentication mechanism. It is expected that this management will be accomplished in-band across the IP overlay from a regional management center (such as an ITSC).
- Contractor network connections to the MAN are described in Chapter 3.7.6. These connections must be consistent with this architecture and should be negotiated through the ITSC.

B.3.3 MAN Specifications

Technical specifications are required to ensure that network services and capabilities are consistent for MANs across the Navy and Marine Corps. In this way, performance is assured and large organizational customers and operators functioning in MANs across the DON see consistency in the following areas:

- ***Interconnectivity to other joint services*** – Seamless connectivity among Army, Navy, Air Force, Marine Corps, and Coast Guard must be ensured.
- ***Virtual private networks (VPN)s*** – Selected commands require VPNs and will require that the ITSC support and manage these VPNs.
- ***Voice, video, imagery, and data over ATM*** – Network-centric warfare requires the full spectrum of telecommunications.
 - ♦ For data, the MAN should support end-to-end SVCs.
 - ♦ For voice, regional planners and designers should ensure that NI-1-compliant ISDN telephone service is available to support the Secure Terminal Equipment (STE) secure telephony.

- ***Routing and addressing*** must be managed and supported for all organizations within the MAN, including participating joint services.
 - ♦ MANs must interface into one or more ATM routing domains, become part of a DON-wide ATM Network Service Access Point (NSAP) addressing plan, and participate in a consistent, hierarchical routing architecture.
 - ♦ Unrestricted IP traffic between campuses (that is, the associated IP networks) is supported on the backbone using the routing scheme described in this document.
 - ♦ In order to minimize the size of internal routing tables, MANs should obtain a Classless Inter Domain Routing (CIDR) block sufficiently large enough to provide IP service to all organizations within the MAN. The required size is based upon the number of campuses and the population of each campus.
- ***Name resolution*** must be managed and supported for interconnectivity throughout the DON enterprise.
- ***Integrated encryption*** – Selected commands will require encryption services for voice, video, and data.
- ***Circuit management for voice, video, and data*** – Selected (if not all) commands will require management (to include provisioning and billing) of their circuits.
- ***Remote management of CANs through the MAN*** – Closed networks will require the ability to pass management traffic through the MAN to other central management centers outside the MAN (such as TIMPO at San Antonio, Texas).
- ***“.com” traffic allowed to the campus*** – Commercial traffic must have the ability to access closed networks within the MAN while maintaining overall information security.

B.4 Service Considerations

Decisions relating to the service provider affect more than best value. In many cases, particularly with ATM technology, the source of the services may determine what is possible to implement. Examples include support of SVC or Private Node-Node Interface (PNNI) hierarchical-based routing that many commercial service providers do not support. Also, ownership of some components may be the sole answer to implementing some desired network capabilities. This section provides a discussion of buying versus leasing, a checklist of services, and a balanced MAN implementation that addresses the needs of all customers.

B.4.1 ATM MAN Ownership

The intended ownership (public versus private) of the components of the MAN is a major factor in the consideration of planning and implementation issues. Ownership carries with it control, including many factors that affect the ultimate level of network services and security. Relevant considerations include Naval requirements, commercial services available, and business case analysis. The ownership answer is by no means the same for every MAN and must be determined based on careful analysis.

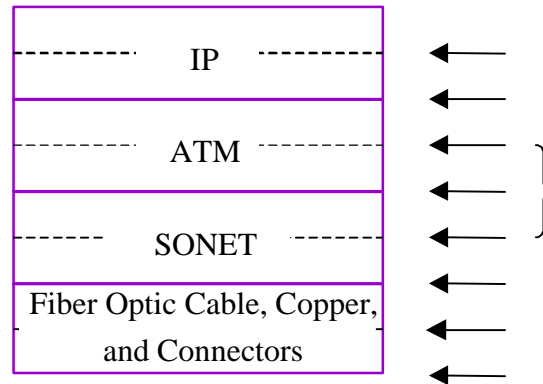


Figure B-7. Network Layers Relevant to MAN Planning

The layers shown in Figure B-7 depict the network components and services that are roughly related to the layers of the Open System Interconnect model. The components and services are discussed in the following bullets:

- IP services will be DON-owned in all envisioned circumstances.
- Naval ownership of ATM services provides many advantages, including performance, security, and reliability. Where justified by mission requirements, existing service offerings, and business case analysis, DON ownership should be an option.
- SONET (including both the cable runs between SONET multiplexing devices and the standards for transporting signals on a fiber optic cable, for multiplexing data, and for frame generation) is usually obtained as a commercial service. In selected instances, the DON may want to own the SONET layer. Examples might include a Naval mission requirement under which bandwidth at the SONET layer is established. Implementation can be made practical by placing a Naval SONET layer on top of the commercial provider's SONET service. In certain cases, direct ATM interconnectivity can occur with an intervening SONET ADM layer (the physical layer will still be SONET but there will be no SONET ADMs).
- Ownership of the MAN physical connectivity (including the fiber optic cable, copper, and connectors) is not supported by Naval requirements or by business case analysis. Physical connectivity will be obtained as a service from commercial service providers, except in special cases such as remote/overseas sites like Guam, Bahrain, and Sigonella.

B.4.2 Selecting a Commercial Service Provider

The selection of a service provider is as important as determining the specific network services to be obtained. In many areas, a single service provider is not able to meet the Naval requirement for services. Many service providers are either inexperienced or unwilling to support the desired protocols or provide the service level needed for the DON MAN. The following ATM areas of support are provided as a guide for evaluating service providers:

- **Service categories** - Tuning the ATM to the particular MAN for voice, video, and data may require a variety of service offerings beyond Variable Bit Rate and Constant Bit Rate. The service offering should support the desired service quality and fit the specific applications of the MAN.

- **Quality of service** - Working with service categories, QoS parameters are a determinant of network performance and must support the desired service quality.
- **Size of connections** - Bandwidth estimates should be slightly higher than the expected use and should be matched to the service category.
- **Support of ATM protocols** - SVCs can be an attractive option in many MANs and should be implemented when appropriate. ATM protocols that support voice and multimedia are similarly important.
- **Network optimization** - Maximizing the throughput of Transmission Control Protocol/Internet Protocol (TCP/IP) networks to match the characteristics of the MAN can be performed by knowledgeable service providers.
- **Network monitoring** - Detecting and isolating faults as they occur should be performed to the level of granularity that permits tasking the appropriate repair technician.
- **Support of LANE and MPOA** - These ATM protocols are complementary and are important to support IP and other legacy protocols. They also introduce issues regarding placement of servers, caching of addresses, and the set-up of VPNs.
- **Traffic management** - The service provider operating capabilities and policies affect performance, particularly when the MAN is supporting service interworking.
- **Over-subscription** - Service providers vary in their tolerance for the bursty nature of certain data applications.
- **Automatic rerouting** - Responses to failed scenarios are not equal among all service providers and must be clearly delineated and understood.
- **Service level agreements** - Specific written agreements should cover relevant aspects of service availability. Recommended areas include allowed downtime per year, mean time to repair, and mean time between failures; service performance such as cell loss; and service installation such as intervals for new ports and PVPs. Effective agreements include a sufficient penalty and means to identify missed service levels.

B.4.3 Balanced Implementation

Determining the balance of enterprise, regional, and organizational functions is a critical success factor in the network connectivity and services implementation. Two overarching ideas (Singularity of Purpose and Autonomous Networks) must be supported.

- **Singularity of Purpose** - Decentralized (i.e. regional and campus) infrastructure performance improvement and consolidation cannot occur without a common DON enterprise network. This assumes that planners incorporate the requirements of all “area” users in the MAN implementation.

The MAN must provide connectivity services to all user groups defined by any Navy and Marine Corps organization at or above the autonomous network layer. Users may be in multiple communities of interest, but at any one time user computers are probably only connected to a single autonomous network.

- **Autonomous Networks** - A number of organization networks are built on top of the enterprise connectivity and bandwidth infrastructure. Autonomous networks, to varying degrees, may be independent of each other from the standpoint of media, technology, security, management, and other characteristics. They are more than virtual networks and may include some physical components that are unique to the particular autonomous network. Autonomous network examples include the fleet intranet and other networks which support the Marine Corps (MCEN), NIPRnet, SIPRnet, BUMED, and SYSCOMS. The fact remains that when possible, these organizations must be part of the DON enterprise network to make information superiority and RMA possible.

To the extent possible, owners of autonomous networks will work with the MAN/regional managers to use the physical components of the MANs to enhance interconnectivity and to reduce duplication and cost.

B.5 Outcome Based Implementations - Metrics

The MAN template places emphasis on design guidance and establishing service levels or outcomes. Clearly, evaluation of alternatives and definition of desired service levels are an integral part of planning and implementation. They are also necessary to evaluate and determine alternative solutions. A basis for optimizing any decision approach should be to view decision-making in the context of the enterprise, region, MAN, and campus and not just in terms of the immediate entity. Functions that should be aggregated under the next level should not be continued at the present level when no performance/cost rationale exists.

A critical success factor in the life cycle of any MAN planning and implementation is the development of the RFP. A government-developed and -prepared RFP should provide measures that furnish a satisfactory representation of Naval requirements.

The suggested guidance for selection of a commercial service provider is a beginning point. These and other relevant requirements factors should be stated in terms of empirical measures that set expectations and provide accountability for performance.

The template will be measured and evaluated under six categories: **security, functionality, interoperability, availability, performance, and cost**. A level of service and in some cases, the specific technologies required, should be described for each. The following are provided as a guide; actual requirements may warrant adjustment of these values and should be based on solid documentation.

B.5.1 Security

A number of security requirements are allocated to the MAN. Cryptographic equipment is required for the protection of classified (and possibly SBU) information. Redundancy should be used. Mechanisms to control access to critical networking components such as switches, multiplexers, and routers should also be used.

- Information confidentiality and integrity
 - ♦ Level of service: must provide sufficient level of protection for information confidentiality and integrity. When a DON-controlled MAN is employed, adequate physical and personnel

security must be established to protect information at the SBU level. When a non-DON-controlled MAN is employed, National Security Agency (NSA)-approved encryption devices shall be used to encrypt data before entering the MAN so that it will be unclassified. However, the service provider must ensure that adequate physical and personnel security is established to protect traffic statistics of DON information traversing the MAN.

- ♦ Technology: KG-75 Fastlane, physical security, and personnel security
- Survivability: specifically relating to security; see also performance
 - ♦ Level of service: must provide protection against denial of service threats (including hostile IW, human error, etc.) commensurate with the criticality of information that will traverse the MAN.
 - ♦ Technology: token-based access control for network components, intrusion detection systems, network management systems, contingency planning

B.5.2 Functionality

The following network infrastructure capabilities that are necessary to effectively and efficiently support the operational mission and requirements must be clearly defined.

- ATM
 - ♦ Level of service: support voice, video, imagery, and data
 - ♦ Technology: end-to-end switched virtual circuit (SVC) and permanent virtual path (PVP)
- MAN connectivity
 - ♦ Level of service: each MAN is provided access to two geographically-separated WAN switches
 - ♦ Technology: N/A
- Switch functionality
 - ♦ Level of service: must support constant bit rate (CBR) QoS
 - ♦ Technology: non-blocking and provide separate queuing for different QoS classes. WAN switches will support at least 2048 switched and/or permanent virtual circuits (PVCs) per interface.
- Availability
 - ♦ Level of service: accessible to all MANs and outlying bases. It is assumed that the MANs will be carrying data for mission critical applications, therefore, an availability of 99.99 percent is an acceptable minimum. (The level of four nines forces redundancy; a single-threaded system (single points of failure) that fails cannot be repaired fast enough to meet this level.)
 - ♦ Technology: N/A

B.5.3 Interoperability

The components of the network infrastructure must interconnect and efficiently and effectively communicate signaling and other information transfer data.

- Service delivery points
 - ♦ Level of service: must support the full suite of Information Technology Standards Guidance (ITSG)-cited ATM protocols and addressing schemas
 - ♦ Technology: N/A

B.5.4 Performance

The MAN service must be sufficiently qualified to meet the following information transfer requirements to support the Naval mission.

- Bandwidth
 - ♦ Level of service: minimum of OC-3
 - ♦ Technology: redundant dual-homed
- Delay
 - ♦ Level of service: maximum end-to-end delay for a CBR service connection should be less than 200 us/switch for processing and queuing plus the necessary propagation delay. The maximum end-to-end cell delay variation should be less than 1 ms.
 - ♦ Technology: N/A
- Latency
 - ♦ Level of service: if switched point-to-point SVC or SVP is supported, the average latency in completing a call setup should be less than 100 ms/hop.
 - ♦ Technology: N/A
- Network availability
 - ♦ Three principles of high availability are established: (1) eliminating single points of failure, (2) reliable crossover, and (3) prompt notification of failures as they occur. All three must be accounted for in performance metrics.
 - ♦ Network monitoring is an issue that requires visibility. A problem in monitoring network availability is determining the failure point when something breaks. The normal situation is for the commercial vendor and base telecommunications to each deny responsibility. Through the ITSC, the capability should exist to immediately sort out the failure point and to call the correct repairman. One means is SNMP visibility of the vendor's network by having access in the ITSC to the real-time availability data that the vendor system monitors are indicating. Having this access should be part of the regional/MAN service provider specification.
 - ♦ Level of service: 99.99 percent
 - ♦ Technology: no switch or physical circuit is a single point of failure

- Survivability (specifically relating to performance; see also Security)
 - ♦ Level of service: vulnerability to forces of nature, acts of humans including enemy action (e.g., backhoe, loss of power, terrorist strike)
 - ♦ Technology: no switch or physical circuit is a single point of failure.
- Quality of Service
 - ♦ Level of service: ability to support a number of service classes based on the traffic type, each with an associated QoS parameter
 - ♦ Technology: N/A
- Mean Time to Repair
 - ♦ Level of service: premium service: less than 2 hours
 - ♦ It is important to note that this decision is highly impacted by the type of system being supported. If the system has double redundancy (is dual-threaded) or triple redundancy, the requirement for premium service may not be justifiable. The trade-off is between adding more redundancy with next day service or a single-threaded system with immediate on-call maintenance.
 - ♦ Technology: N/A

B.5.5 Cost

Implementation cost is a significant category for judging value and must be driven by the context of the level of service.

B.6 Evaluating MAN Products and Services

The selection of MAN products and services should be made by an organizational team comprised of representative operational and IT technical experts.

The determination will be made based on a balanced scorecard approach using as a basis the characteristics defined in Section B.5. Addressing the proposals of each competing product or service should be in accordance with the criteria in Figure B-8.

| Network Characteristic | Raw Score | Weighting (.xx) | Adjusted Score |
|------------------------|-----------|-----------------|----------------|
| Network Security | X | .30 | X |
| Functionality | X | .20 | X |
| Interoperability | X | .20 | X |
| Performance | X | .30 | X |
| Total | — | 1.0 | X |
| Total Merit / Cost | — | — | X/\$YY |

Figure B-8. Product / Service Evaluation

1. Four of the five characteristics will be scored based on a numerical rating system (1-10). The supporting subsets of each characteristic will be evaluated and aggregated to determine the value of the characteristic.
2. Also, each characteristic will have a weight factor based on its judged importance to the overall region/MAN/CAN mission. The total of these individual weight factors will be 1.0. Recommended weight factors are provided.
3. The numerical value for each characteristic is multiplied by its weight factor and the total of the five adjusted scores equals the relative merit for the source provider.
4. Cost will be viewed as a separate variable. The relative merit will be expressed in a total merit to cost ratio (e.g, X / \$ YY).

The team should complete an analysis of the alternative providers and present a recommendation to the regional commander/cognizant decision-making body. Included in the presentation will be a recommendation for the administration and funding of the product or service.

This page intentionally left blank.

Volume I, Appendix C - Table of Contents

| | |
|---|------------|
| C. Campus Area Network Design Template..... | C-1 |
| C.1 Purpose | C-1 |
| C.2 Drivers | C-2 |
| C.3 Elements/Features/Specifications | C-3 |
| C.3.1 Topology | C-3 |
| C.3.2 Physical Overlay..... | C-6 |
| C.3.3 Switches, Routers, and Hubs | C-8 |
| C.3.4 Transmission Overlay | C-10 |
| C.3.5 ATM Technology Overlay | C-11 |
| C.3.6 IP Overlay | C-14 |
| C.3.7 Switching and Routing Overlay..... | C-17 |
| C.3.8 Voice Overlay..... | C-18 |
| C.4 Functional and Performance Specifications | C-19 |
| C.4.1 Security | C-19 |
| C.4.2 Functionality..... | C-20 |
| C.4.3 Interoperability | C-21 |
| C.4.4 Performance | C-22 |
| C.4.5 Cost..... | C-23 |
| C.5 Evaluating CAN Products and Services | C-23 |

This page intentionally left blank.

C.Campus Area Network Design Template

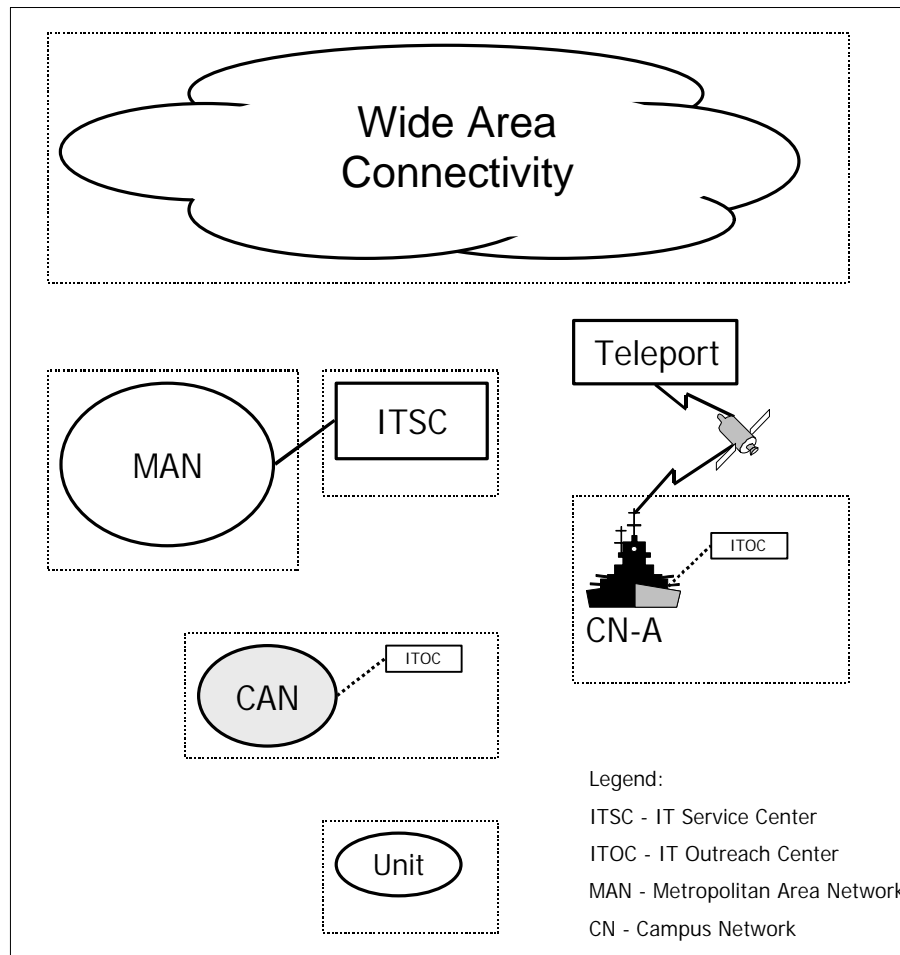


Figure C-1. High Level Components of the TI Architecture

C.1 Purpose

Figure C-1 depicts the relationship of the campus area network (CAN) to the other components of the ITI architecture. Navy and Marine Corps campuses vary widely in their network connectivity requirements but they must all function together effectively within the DON enterprise network.

The purpose of this CAN template is to establish a set of general guidelines within which a campus network may be built. The template provides a consistent means to address campus network components, identifies appropriate instances for specific ITI solutions, and details a list of implementation considerations. The guidance is presented in a manner that allows considerable variance to address local needs, to support a best value solution, and to align campuses within the context of a DON enterprise infrastructure solution.

The architecture guidance for shipboard ITI is included in the CAN template. The majority of the campus architecture guidance is common to both terrestrial and shipboard environments. In the absence of shipboard-specific guidance to the contrary, the general campus guidance should be

implemented for shipboard networks. When there are shipboard unique architecture requirements, they will be noted in blue italics.

C.2 Drivers

The campus contains the tactical and support units that perform the Navy and Marine Corps mission. Performance of the mission is directly affected by the quality of the information provided to the units operating on the campuses. The drivers that determine quality are as follows:

- **Interconnectivity.** A robust IT infrastructure is rooted in the interconnectivity of the networks and systems that comprise the DON enterprise information system. It is a fundamental element of the RMA. Global connectivity through the WANs, MANs, and CANs must be supported by a well-defined and supported ITI architecture.
- **Interoperability.** Implementation of campus area networks must provide for interoperability among units within campuses, among campuses, between regions, and with external organizations. *Shipboard networks must provide for interoperability among units within ships, between ships, and with external organizations.*
- **Seamlessness.** Infrastructure systems at the campuses must be able to identify and communicate with systems across the DON enterprise infrastructure, without the need for special human intervention, on other campuses *or shipboard networks*, in regions, and outside the DON enterprise.
- **Integrated, Supportable Solution.** The campus architecture must support integration at both the campus *or shipboard network* and MAN/region levels. Integration of voice, video, and data is required to reduce infrastructure duplication and support cost. The integration of network services across the MAN must reduce duplication of functions and enable supported organizational units to focus on their primary missions. Regional supportability – especially the ability to remotely manage devices at the ITSC – is an important component.
- **Affordability.** The campus networks and services must be realigned to leverage resources, eliminate redundancy, and consolidate underused or common functions to enable the same or better network services at reduced cost.
- **Availability.** *Unlike many of the functions supported by CAN or MAN networks, shipboard networks support tactical functions that are critical to personnel safety and mission success. Shipboard networks will require higher levels of redundancy and a more robust design philosophy than their shore-based counterparts.*
- **Latency.** *Shipboard networks will have to support real-time applications with latency requirements in the millisecond range.*
- **Mobility.** *Shipboard networks move from homeport to sea to foreign ports and require connectivity into the global ITI under each of these circumstances.*
- **Bandwidth disadvantages.** *Shipboard networks will never have the unlimited “uplink” bandwidth potential that is available to shore-based campus networks.*

C.3 Elements/Features/Specifications

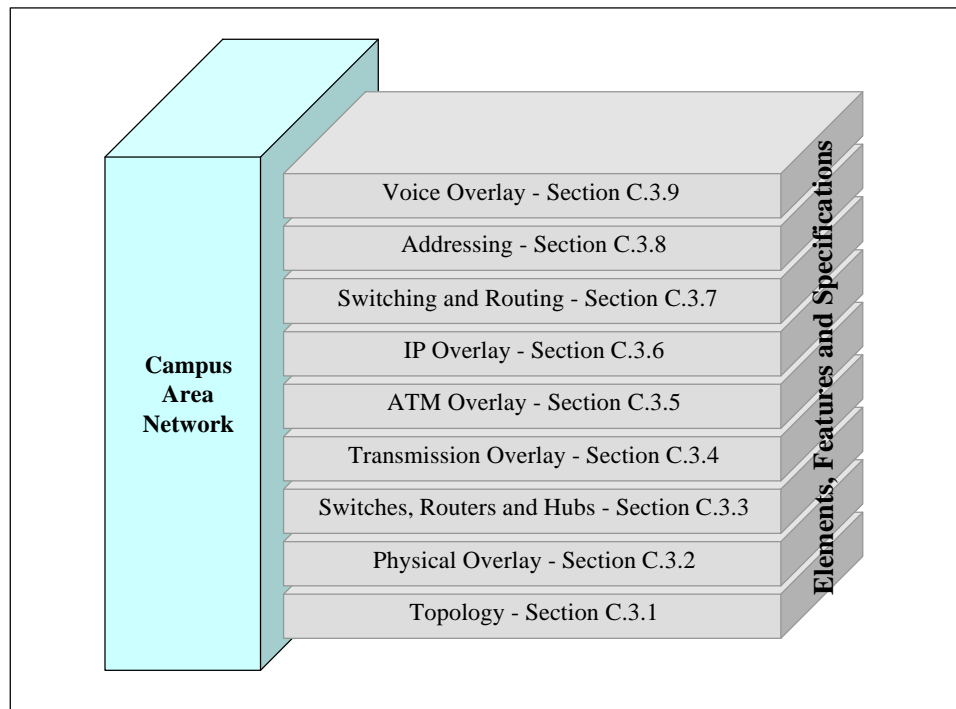


Figure C-2. Layered Description of the Campus Area Network

The complexity of a campus area network is illustrated by the elements depicted in Figure C-2. These layers roughly correspond to the Open System Interconnect (OSI) protocol stack. Each of these is discussed in the nine sub-sections that follow, including a description of important features, recommended implementations, and required specifications.

C.3.1 Topology

The elements of the CAN are depicted in a notional drawing shown in Figure C-3. To summarize, the campus is connected to other campuses through an ATM Private Network-Network Interface (PNNI) service or Permanent Virtual Path (PVP) mesh, normally OC-3 or higher, as provided by the MAN. The campus premise switch is the demarcation point for the campus network infrastructure.

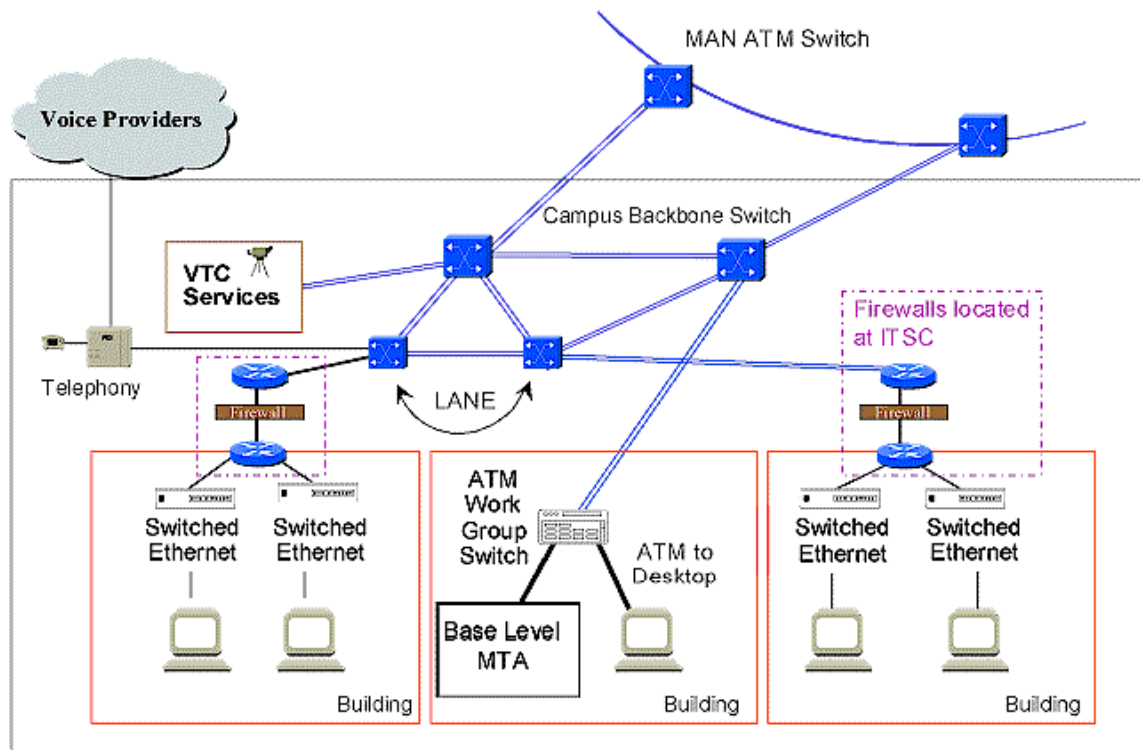


Figure C-3. Notional Campus Area Network

Within the campus, the information architecture supports a multi-media infrastructure through a heterogeneous ATM/IP technology solution. The voice infrastructure will eventually be integrated into the multi-media infrastructure using ATM technology. The CAN provides interconnectivity to a node at each of the buildings on the campus. The topology of the multi-media infrastructure is dependent upon the physical layout of the campus, traffic demand patterns, the degree of required survivability, and other location-specific considerations. Basic guidance is provided to address these expected implementation differences.

The multimedia transport topology at a particular campus is defined in the campus blueprint and should be consistent with the general design and implementation guidance provided in this template.

The elements of shipboard networks are depicted in Figure C-4. To summarize, the shipboard network is connected to shore-based or other shipboard networks via satellite or through a pier-side cable connection. An ADNS router is the demarcation point for the shipboard network infrastructure.

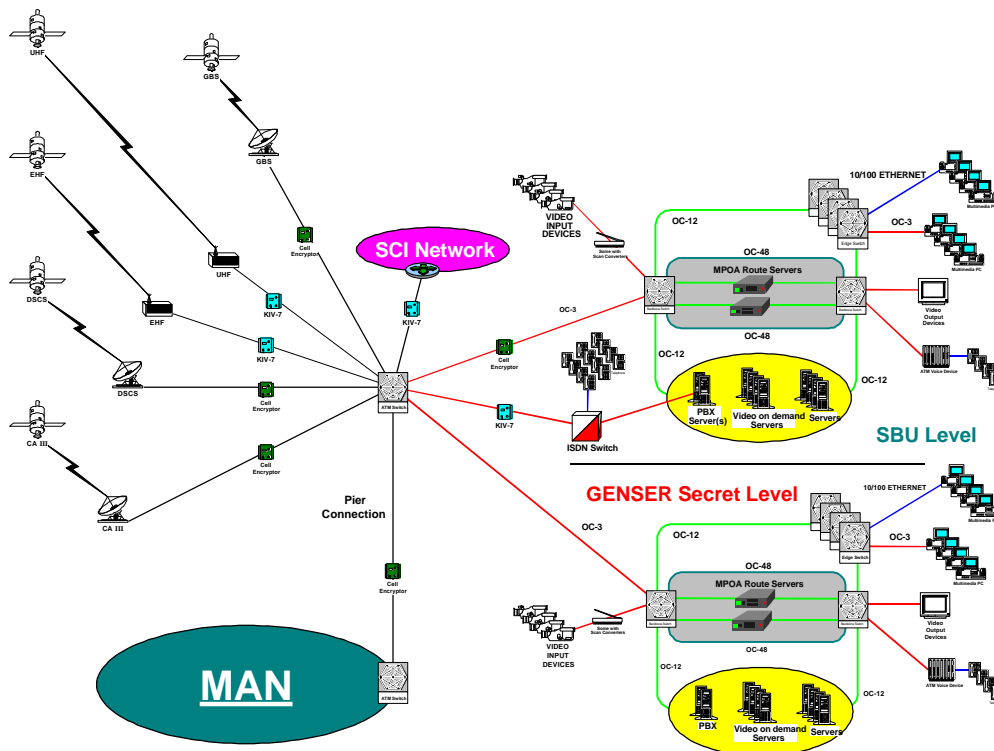


Figure C-4. Notional Shipboard Network

Shipboard networks provide interconnectivity to sensors, processors, weapons systems, and user workstations throughout the ship. The multimedia transport topology at a particular shipboard network is defined in the shipboard network blueprint and should be consistent with the general design and implementation guidance provided in this template.

C.3.1.1 Discussion of Alternatives

Three common methods can be used to connect networks.

- **Ring:** Fiber Data Distribution Interface (FDDI) is a common example of a ring topology. FDDI is a high performance fiber optic token ring technology in which the “token” is passed from one node to another until it reaches the desired node. Other network technologies such as SONET can be configured as a ring in which each circuit is given a provisioned time slot.
- **Star:** A star network connects each device to a center location by a point-to-point link. The center device is usually called a hub or concentrator in which the center point may be passive, active, or intelligent. This topology can be extended by connecting one star to another. Networks based on ATM and switched Ethernet are traditionally configured as stars.
- **Mesh:** In a mesh network, each component is directly connected to several other devices in the network. A full mesh indicates that each device is connected to all others while a partial mesh indicates that not all links are in place.

There are two logical alternatives to the recommended shipboard network architecture:

- *Multiple “stovepipe” networks interconnected via routers. This architecture, which currently exists on a number of ships, does not support the desired levels of interoperability.*
- *A single multi-level security (MLS) network instead of the separate Genser Secret and SBU networks. This architecture is preferable but is not practically realizable today with existing security components.*

C.3.1.2 Recommended Topology Implementation

Campus networks are typically built with a backbone that can have its capacity increased without affecting end users (by swapping out hubs and switches). Backbones may consist (initially) of 10Base-T Ethernet, 100Base-T Ethernet (fast Ethernet), 1000Base-T Ethernet (gigabit Ethernet), FDDI rings, or ATM meshes. The border of the network -- the interface a user sees -- is most commonly 10Base-T Ethernet. For those users with higher capacity requirements, the backbone network can be extended to them.

Shipboard backbones can have their capacity increased by swapping out edge devices and switches. Backbones should consist exclusively of fiber ATM mesh networks. The border of the network may consist of copper 10Base-T, fiber 10Base-FL, copper or fiber Fast Ethernet, or fiber OC-3 ATM. Gigabit Ethernet and 25 Mbps ATM should be avoided for shipboard network users.

The network topology implementation should support future growth and accommodate future networking technology. Cost of both fiber optic and copper media are very small compared to the labor cost of installation. Therefore, it is almost always economical to put extra fibers in between-building runs (known as ‘dark fiber’) and extra Category 5 copper runs within buildings that can be brought into the network later. Campus backbones (the runs between buildings) should normally have alternate routes (dual rings or mesh topologies) so that the effect of single points of failure can be contained.

To maintain the flexibility to support future growth and accommodate future networking technology, sufficient fiber optic media must be installed in ships in as adaptable a configuration as possible. To achieve this objective, the physical topology of the fiber optic cable plant between core nodes will, in general, be a partial mesh topology with dual-homed connections to the edge devices.

Each core node will be connected to at least two other core nodes via separate, physically diverse paths to enhance survivability of the backbone transport system. This topology allows implementation of physical and logical ring, point-to-point, hub, or mesh topologies which can be used in an ATM or combined ATM and SONET transport environment.

C.3.2 Physical Overlay

Cable and cable equipment must meet current and future requirements for data transmission, electrical characteristics, and topology. Fortunately, manufacturers have boosted data transfer rates on copper twisted-pair wire so that it can meet some bandwidth demands to the desktop. *Copper cable should not be used for the backbone of shipboard networks.* The advantages of fiber cable, however, are many and include support of many times higher transmission rates and far greater security because they produce no emissions and are not affected by radiation.

C.3.2.1 Discussion of Alternatives

This discussion of alternatives includes both metal wire and fiber-optic cable. They are addressed in the context of new installation/replacement media and retention of existing media.

C.3.2.1.1 New installation/Replacement

To maintain flexibility, support future growth, and accommodate future networking technology, sufficient fiber optic media should be installed on the campus in an adaptable, standards-based configuration. The exact number and type of fibers to be installed on the campus backbone should be based on careful analysis of current and future growth. A rich mixture of 62.5-micron multi-mode and 8 micron single mode fiber is recommended.

Although fiber is the preferred medium for outside plant (between buildings), it may be permissible to install new copper in certain situations such as within a building or long runs of low-bandwidth signals (such as remote guard shacks). Category 5 (at least) Unshielded Twisted Pair (UTP) is the recommended copper media. (Category 5 cable is suitable for POTS telephone use as well as LANs. But Category 3 telephone cable is unsuitable for LAN use. In the interests of a single cable plant, Category 5 is recommended for both applications).

For shipboard low-end workstations that do not need the multimedia capability provided by fiber interfaces, standard Screened Twisted Pair (SCTP) Category 5 is the recommended copper media.

Coaxial cable should not be installed on the campus for inside or outside plant connectivity unless there are special circumstances that dictate its use. *Coaxial cable should not be installed for shipboard networks.*

C.3.2.1.2 Retention of Existing Media

Although fiber is the preferred medium, metallic (copper) transmission media should be retained or reused where the evaluation of performance, mission priority, security, and cost dictate. Asymmetric Digital Subscriber Line (ADSL) is a technology that may be effectively and economically used to salvage existing telephone wiring in a campus environment. ADSL can provide point-to-point capability with capacities in the range of T1.

C.3.2.2 Recommended Physical Implementation

The notional CAN in Figure C-3 shows two ATM top level switches for connection to the MAN and, ideally, this concept would be repeated for connection of these top level switches to all building switch/routers.

Fiber links should be used to connect floors of buildings to the locations of the interface points to the ATM switches. *Fiber links should be used for all shipboard backbone networks.*

The exact number and type of fibers to be installed on backbone routes between network nodes, devices, and buildings will be based on engineering analysis of current and future growth requirements. In most cases, the cost of installation far exceeds the media cost, so overbuilding is usually the right answer to the analysis.

The decision to use multi-mode or single-mode fiber to connect core switches should be based on current and near-term bandwidth needs and cost. For example, multi-mode fiber is cost-effective for OC-3 long distances, but single mode fiber is needed for all OC-12 and OC-48 connections. In all cases, single-mode fiber should be run between core switches and between buildings. In FDDI

(100Mbps capacity), multi-mode fiber reaches around 2 km between nodes, and single mode fiber will reach 20-30 km.

From the building switch/routers to the server and desktop sites, several alternatives can be used. For maximum reuse of existing cable plant and for least-cost network interface cards (NICs), standard unshielded twisted pair (UTP Category 5) copper cable can be used. *Again, SCTP has been determined to be clearly superior to UTP for both noise immunity and security and should replace UTP where appropriate.*

For very high-end servers such as multimedia servers or high-end workstations (bandwidth requirements in excess of 100 Mbps), fiber to the office or location can be used. The CAN implementation does not mandate exclusive use of fiber, but should evolve to fiber where it can be justified by cost benefit analysis.

Given future bandwidth needs, an initial investment in a robust set of multi-mode and single-mode fiber will be justified.

The EIA standards for Category 5 installations include documentation standards. Since overbuilds are recommended, documenting the dark fiber and unused UTP is important so that it can be located later for use. Therefore, a configuration management scheme that documents the as-built configuration is important.

C.3.3 Switches, Routers, and Hubs

Switches, routers, and hubs are the major components of LANs, and their selection requires careful attention. Generally, campus implementers are wise to obtain specific components from a single vendor for interoperability and support reasons. Also, these components should have management agents embedded in them to support remote electronic monitoring by the Information Technology Service Center (ITSC).

Switches

FDDI and all variants of Ethernet are standardized on the IEEE Committee 802 model and use a common (802.2) logical link controller. In practice, this means that all these LAN technologies can be bridged together with appropriately-configured switching hubs.

Interoperability of switches is based on standards-based routing, and support for dynamic routing is accomplished through the PNNI protocol. Prior to selecting a vendor, have a thorough interoperability test that has been demonstrated on networks of similar size. The burden can, and should, be placed on the vendor. A suggested contracting option is to specify that a detailed test plan must be demonstrated before final award. Should the test plan be extensive, a preliminary award can be made to the potential vendor, but the final award should be held until the successful passing of the test plan.

The sizing of core and secondary switches should be based on current and near-term cost-effective technologies. For example, OC-12 is now a reasonably-priced interconnection link between core switches. While the current population/use may not fully justify a switch initially, expected growth should be accommodated by the design of the switch chassis.

Routers

Two classes of routers can be implemented in DON networks. Conventional routers provide network-layer route computation and packet forwarding in a single physical device. Virtual routers are comprised of route servers that are used to perform route calculations and non-routing edge devices.

The two major classes of routing information protocols currently implemented in contemporary networks are distance vector (Bellman-Ford algorithm) and link state. The Open Shortest Path First (OSPF) link state routing protocol is the most predominant routing protocol on DON networks and is the recommended interior gateway protocol.

Routers with ATM interfaces will be needed for routing between the switched ATM backbone network and existing LANs. These routers should be capable of supporting a minimum of OSPF plus any locally-required routing information protocols (consistent with the ITSG). If virtual LANs (VLANs) are implemented, these routers must also be able to perform routing between multiple VLANs on various network segment types (such as Ethernet, Token Ring, and ATM). Future versions of edge switches may also support multi-protocol routing, which reduces the necessity of a physical router. The ATM router, when used as connection between the ATM CAN and the existing legacy LANs, should be capable of supporting the following features:

- LAN Emulation (LANE)
- Classic IP and Address Resolution Protocol (ARP) over ATM as defined in RFC 1577
- Multi-protocol encapsulation as defined in RFC 1483
- PVC and SVC connections
- ATM Forum UNI signaling (UNI 3.0/3.1, UNI 4.0)
- AAL3/4, AAL5
- Multi-protocol routing over ATM (MPOA) (routing of IPX, DECnet, and Appletalk are local options -- extra-campus routing is IP only).

For campuses that run ATM and are supported by a MAN that uses PNNI, the CAN template assumes that the ATM backbone switches on the campus also perform routing using PNNI.

Hubs

Hubs fan the circuits out to the borders of the LAN. They also perform bridging between separate, disparate LAN technologies. Furthermore, if configured in accordance with DON ITI architecture recommendations, the hub can provide significant network monitoring data for centralized monitoring and management in the ITSC. The following guidance applies to hubs.

Protocols Supported

- Ethernet. The vast majority of end systems (such as desktop computers) have 10Base-T (Ethernet) (IEEE 802.3) interfaces. Hubs should therefore have 10Base-T ports by default.
- Ethernet and Fast Ethernet. Hubs that can handle both 10Base-T and 100Base-T (fast Ethernet) are quite common. These hubs support two functions: attach a mix of 10Base-T and 100Base-T end systems to the hub and provide 10Base-T for the end systems and 100Base-T for the campus backbone. Hubs can also bridge the varieties of 802.3 Ethernet, FDDI, and token ring with the appropriate interfaces.

Types of Hubs

- Shared media hubs are at the lower end of the price range and have the disadvantage that all end users receive all traffic on the LAN whether it is addressed to them or not. *Shared media hubs should not be used on ships.*
- Switching hubs have the identical external interfaces of shared media hubs. But switching hubs provide each end system with a bridging function that eliminates contention on the connection between the hub and end system—the hub sends traffic out one port that is only destined to the end system attached to that port. One means of controlling congestion is replacement of a shared media hub with a switching hub.

Network Monitoring

Hubs that contain an SNMP agent are referred to as intelligent hubs. Once programmed, the SNMP agent interacts with the ITSC network operations center so that the network monitoring software at the ITSC can determine remotely what is and is not working and the levels of congestion at the network hub. (Additional information is contained in Appendix D under Performance Management).

The cost of the SNMP agent is a marginal increase over the cost of an unmanaged hub; hubs should be purchased with SNMP agents installed.

C.3.4 Transmission Overlay

The SONET protocol has replaced much of the world's PDH equipment (Plesiochronous Digital Hierarchy). All future systems should support the SONET protocol, be it circuit emulation via ATM or SONET Add/Drop Multiplexers (ADMs). PDH circuits (DS-0, DS-1 and DS-3) can ride on this equipment.

C.3.4.1 Discussion of Alternatives

DS-1 (formerly T-1) is basically a conditioned telephone line and is the most common digital leased-line service and provides 1.544 Mbits/sec. DS-3 (formerly T-3) is also common and provides 45 Mbps.

SONET defines a fiber-optic transmission that is standard for cell relay. Typical line speeds are OC-3 (155 Mbps) and OC-12 (622 Mbps). When appropriate, SONET provides a number of advantages:

- Higher capacity, scaleable bandwidth from OC-1 (51.84 Mbps) to OC-192 (9.6 Gbps)
- Improved Survivability - protection switching in less than 50 msec
- Bandwidth management allows a customer to obtain more, less, or redirected bandwidth as needed
- Rapid provisioning - ability to provision a service in minutes as opposed to days or weeks as is typical in the existing environment
- Improved operation, maintenance, and provisioning - voice channels for operators and data channels for network management

- Single channel visibility - frame interleaving enables single channel visibility via SONET ADMs. Depending on the equipment used, individual channels down to the DS-1 or DS-0 level may be added or dropped.

C.3.4.2 Recommended Transmission Implementation

The preferred choice for transmission systems is an ATM switch fabric with optional SONET ADMs. The use of an ATM-only switch fabric has cost benefits and allows for dynamic circuit creation and interconnection to the WAN. The added expense of SONET is warranted for certain applications such as real-time data acquisition *and circuit-switched voice*.

When deploying ATM in the core backbone, equipment from a single vendor is desired because of the state of the protocol implementation (ATM protocols are still evolving and not all features are fully implemented by all vendors).

At campuses where SONET is implemented, SONET equipment from a single vendor should be installed to facilitate interoperability; simplify maintenance and circuit restoration actions; and allow operation, administration, and network management functions to be performed from a single integrated hardware/software platform.

C.3.5 ATM Technology Overlay

ATM is a communications networking technology that supports many types of traffic, including voice, data, real-time video, and imaging. ATM is a cell relay technology, implying that the data packets have a fixed size. Same-size cells provide a way of predicting and guaranteeing bandwidth for applications that need it. The ATM fast VLSI cell switching fabric also provides a means to replace the time-consuming CPU-bound packet forwarding of traditional routers with switches that reduce traffic delays. Unlike Ethernet, FDDI, and token ring that used a shared media, ATM provides an any-to-any connection and allows any pair of communicating nodes to transmit simultaneously.

C.3.5.1 Discussion of Alternatives

The alternatives for the campus are ATM cell switching and IP packet routing coupled with switching hubs; however, no matter what technology is selected for the campus backbone, IP to the host must be supported. Both ATM switching and IP routing have distinct advantages and should be carefully evaluated in relation to the specific campus requirements. Because IP has such a tremendous legacy implementation, any acceptable ATM solution must fully support (and enhance) the IP environment. This section is focused on ATM and necessarily includes IP; additional IP discussion is provided in Section C.3.6. Other supporting architecture information for both IP and ATM is provided in Chapter 2 of this document. The ATM alternatives are addressed in the following context.

- ATM networking solution
- Support for IP (LANE, MPOA)
- Specific ATM Routing Performance
- Security of ATM

C.3.5.2 ATM Networking Solution

There are two potential topologies for implementing ATM: (1) ATM switches in a partial mesh, and (2) ATM with a SONET ring using ADM. Both topologies may also implement VLANs.

An ATM switch is located at the primary node that is the campus ingress/egress for all multimedia transport. Secondary nodes and building nodes, typically implemented as IP routers, may be implemented as ATM switches for campuses hosting organizations requiring ATM access to the DON enterprise network.

Alternative topologies are appropriate in order to accommodate campus geographical features and existing outside cable plant infrastructure. To enhance survivability of the backbone transport system, each secondary node should be connected to at least two other secondary nodes via separate and physically diverse paths.

To support integrated services networks at most DON locations, the recommended implementation is a switched ATM backbone network configured in a mesh topology initially supporting OC-3 (155 Mbps) and later supporting OC-12 (622 Mbps) transmission rates. The core ATM mesh network is designed by connecting each backbone ATM switch to at least two or more other backbone ATM switches.

Connection of legacy LAN technologies (such as Ethernet and Token Ring) to the ATM backbone will be made using LAN switches referred to as ATM edge switches. Redundancy in the switched network base backbone is provided through multiple connections from the ATM edge switches and ATM workgroup switches to ATM base backbone switches. These connections should be capable of supporting 155 Mbps transmission rates. ATM switches should be capable of automatic (transparent-to-end systems) re-establishment of virtual circuits in the event of switch or link failure. This reconfiguration is automatic in routers and available in ATM switches – but the planner must ensure that the capability is received.

C.3.5.2.1 ATM Addressing on the Campus

In instances in which the campus has implemented an ATM network infrastructure, an addressing solution described in Chapter 4 (Section 4.4) is appropriate.

The MAN or region ITSC will obtain an ATM address block to support the participating campuses. Within a MAN/region hierarchy, each campus will be reflected (and announced to the MAN) by a single prefix and mask. Each campus will be allocated sufficient address space.

The ATM address form is based on Network Service Access Point (NSAP) addressing, which means that the end station address includes the host routing information. This includes support of the PNNI routing in which the MAN implements a PNNI routing hierarchy.

For the fleet autonomous ATM net, the DISA Globally Unique Identifier (GUI) or home port geographic address space as offered for deployed forces does not meet DON requirements and will not be used. The Naval solution, under which each ship will have a single prefix and mask, will be developed and presented in a subsequent version of this architecture.

C.3.5.2.2 ATM Routing on the Campus

When a MAN has implemented a PNNI routing hierarchy, the campus should participate as is appropriate.

Transporting IP packets by encapsulation into ATM cells provides significant potential for improved routing performance via LANE and MPOA, as discussed in the following section.

C.3.5.2.3 Support for IP (LANE, MPOA)

The campus ATM solution must support interoperability between devices on ATM LANs and the traditional LANs (Ethernet, Token Ring, and FDDI). It must also support configurations connecting traditional LANs over an ATM backbone.

In the near term, LANE provides a solution for interoperability of LANs independently of any higher networking layer (such as IP, IPX, or DECnet) that uses the services of the LAN.

In the longer term, MPOA is the preferred solution and allows end devices which are attached to separate ATM networks to communicate directly with each other instead of through intervening routers, even when end devices are located on two different subnets (known as cut-through routing). MPOA substantially enhances performance by reducing routing delays. MPOA is defined by an ATM Forum protocol but is not yet fully implemented by vendors and must be implemented in a Switched Virtual Circuit (SVC) environment.

C.3.5.2.3.1 LAN Emulation (LANE)

ATM and IP attached LANs and accompanying end stations must both communicate over the ATM backbone. LANE provides the technology to enable this mixed environment to communicate. This is made possible by having the ATM network "emulate" the characteristics of broadcast LANs and was fully described in Chapter 2.

At the campus, it is possible to implement the LANE functions in an ATM switch, but this should be avoided for performance reasons. While it is possible to provide support for legacy networks (e.g., Ethernet) via LANE services in the core switches, most LANE should be implemented in the edge or intermediate switches.

(For clarity, core switches are dedicated switching systems without additional processing requirements. Edge switches are located at the logical "end" of an ATM network and are typically an ATM switch that supports legacy networks via LANE services or ATM end system workstation/hosts. Intermediate switches lie between the backbone and an edge switch.)

For smaller campus networks, an intermediate switch may not be necessary; the edge switch can connect directly to the core. The decision on when to deploy LANE devices depends on the particular ATM implementation and on how the individual ATM devices handle the LANE processes.

For campus networks which include many autonomous networks, the LAN Emulation Server/Broadcast and Unknown Server (LES/BUS) services should be implemented on the edge switch close to the organizations that connect to the various emulated LANs (e-LANs).

Every e-LAN functions independently using its own LES/BUS and interconnecting multiple emulated LANs which require a router. Instead of routing between multiple physical interfaces like a typical router, the router simply routes the Layer 3 protocols onto multiple Emulated LANs connected through the same interface. The router can be a card that is inserted into an ATM switch or it can be a stand-alone router.

C.3.5.2.3.2 Multiprotocol over ATM (MPOA)

MPOA splits the traditional role of the router into two roles by off-loading packet forwarding from the router to the hosts and edge devices. The ability to separate packet forwarding from other router functions allows a more flexible implementation of the two components. Enabling packet forwarding in a separate device from the router is known as cut-through routing, which uses the Next Hop Resolution Protocol (NHRP). NHRP defines the methods for routers to communicate among each other to determine unknown IP-to-ATM address mappings regardless of IP subnets of the end devices.

Edge devices and hosts discover the ATM address associated with an IP destination by sending NHRP queries to their router servers. The router servers communicate among themselves to learn addresses, and the edge devices can then establish a cut-through route to the destination.

Of LANE and MPOA, MPOA is the preferred solution, but only when planners and implementers are confident that adequate service provider support exists.

C.3.5.2.4 Security of ATM

The following mechanisms are required to secure the DON ATM overlay at the CAN level.

- Fastlane ATM encryption devices are required for connections between secret enclaves. These devices are installed between secret buildings and SBU CANs. (An entire ATM CAN is operated at the SBU system high level.) Enclaves operating at various classification levels can likewise make use of the DON ATM overlay by using key management to keep the classification levels separate. CoIs requiring complete information isolation between themselves and the rest of the DON ATM overlay can likewise be supported using Fastlanes.
- The ATM CAN must be constructed in a redundant fashion such that the failure of a single network component or interconnection will not lead to the catastrophic failure of the CAN. Redundancy of individual building connections to the CAN should be analyzed on a case-by-case basis by using the net subscriber value (NSV) process. See the DON ITSG Section 10.4.3.2.1 for NSV details.
- Bandwidth allocation management must be provided within the ATM overlay. This includes providing authorized administrators with the capability to identify the priority of data transport requests and allocating network bandwidth to the highest priority requests when contention occurs. In addition, the ATM overlay should be designed with the ability to “order and add” additional bandwidth as required by adding additional links to the ATM mesh.
- The ATM overlay must provide mechanisms to ensure that DON-controlled components of the overlay can only be managed by authorized administrators, that they are resistant to penetration attempts, and that they are resistant to ATM signaling-based denial of service (DoS) attacks. To reduce the potential for successful penetrations originating from inside the DON-controlled portions of the overlay, network components that are remotely managed must feature non-spoofable authentication mechanisms. It is expected that this management will be accomplished in-band across the IP overlay from a regional management center (e.g., an ITSC).

C.3.6 IP Overlay

The campus network template assumes that every campus will support end-to-end IP connectivity to every desktop. Within a given campus, a number of autonomous networks are assumed to exist, and this architecture must support routing from one to another. The campus will have full

access to the global Internet; the architecture template will support traffic to the external world. (This will normally occur through the MAN.) The campus network poses security challenges, which are addressed by the interleaving security mechanisms at every segment of the template and providing the gateway to the Internet solely at the ITSCs, from which access can be appropriately managed.

C.3.6.1 IP Addressing for Campuses

DON IP address management includes the efficient allocation and use of scarce IP addresses for afloat, ashore, tactical, and non-tactical networking environments.

Each campus will have a single adequately-sized address space allocated from a Classless Inter-Domain Routing (CIDR) block that is administered by the ITSC. The size of the block should be large enough to support the long-term needs of the campus by providing the ability to add hosts as requirements change. This allocation size will be carefully managed to provide maximum use of available address space. Campuses within the DON enterprise should be supported by the ITSCs for IP network planning, internal routing architecture, and implementing CIDR subnetting.

Each CAN will be subnetted under a single net address/network mask, and the campuses in a particular MAN will aggregate to a single MAN address/network mask. This IP network address management will apply to all IP devices and IP addressing services. IP addresses are officially tracked and assigned only to the level of a Class C address.

The use of subnet masking to subdivide IP address space is supported, but is typically a matter for local shipboard and shore-based network administrators. However, due to the inherent complexities of various IP subnet masking techniques, coordination with the MAN IP service provider is encouraged to ensure the most efficient use of assigned IP address space.

When appropriate, CANs may employ address hiding techniques using the private address space (RFC 1918) for internal connectivity. In such cases, a Network Address Translator (NAT) provides access to the DON enterprise IP network, the Internet, and other external networks. It is useful when a campus uses IP addresses that do not need to be advertised to the DON enterprise network or the Internet.

Dynamic Host Configuration Protocol (DHCP) is used on CANs for IP address management to further ensure the efficient use of assigned address space and to reduce the amount of network administrator time. The latter is accomplished by managing a single DHCP server in lieu of tending to individual end users' workstations. DHCP is limited to standard workstations and other devices whose IP address can be changed without DNS reconfiguration. While DHCP is not specifically intended to support mobile users, its use in this architecture simplifies laptop network attachment at remote locations because each campus establishes a DHCP server. Laptops that are configured as DHCP clients are simply connected to the network and the IP configuration information is automatically obtained. DHCP supports several features used within CANs of the DON enterprise network:

- Automatic assignment of IP addresses and configuration data such as net masks and Domain Name System (DNS) servers
- Automatic reclamation of unused IP addresses resulting in economy of IP addresses
- Centralized administration and management of the IP address space without custom desktop IP configuration

CANs (for the both sea and shore commands) can obtain registered IP addresses from the Naval Computer and Telecommunications Station (NCTS) in Pensacola, Florida, at (COMM) 850-452-3501, (DSN) 922-3501. Addresses can also be obtained on-line at the Navy IP Network Number Registration page at <http://www.netreg.navy.mil/>. This should be done through the cognizant ITSC.

C.3.6.2 DNS Service for IP

Domain name service is a critical support service for IP that matches end system names (host names, Web URLs) with IP addresses. This subject is addressed in depth in Chapter 4. A CAN has three options for DNS service:

1. Rely on the DNS server at the ITSC. This requires that new campus IP address assignments must be communicated to the ITSC's DNS administrator. All DNS requests result in overhead traffic between end systems and the ITSC.
2. Implement a DNS server locally. The local DNS server is the authoritative server for all campus-based end systems and points to the ITSC DNS server as its upstream reference. The disadvantage of this arrangement is that the campus must administer a server, which nullifies the economy of ITSC centralization.
3. Implement a caching DNS server locally. The ITSC server is the authoritative server as in option 1, but the caching function cuts down on overhead traffic outside the campus network. When properly installed, this server is physically resident on the campus network but is managed by the ITSC.

C.3.6.3 IP Routing Services for Campuses

Architecture guidance for routing in the CAN environment depends on two inter-related variables:

- Interior Gateway Protocols (IGP) versus Exterior Gateway Protocols (EGP)
 - IGP** - Routing inside a campus network assumes all users are self-contained, not separated, and contiguous. The preferred IGP is Open Shortest Path First (OSPF) for overall routing.
 - EGP** - The network solution for EGP depends on whether the campus is integrated with the MAN or is independent.
- Campuses that have network services integrated with a regional MAN and ITSC versus those that operate independently as a single Routing Domain
 - Integrated** - Campus networks that receive service from a regional ITSC as part of a MAN will typically be part of that autonomous system, and consequently, OSPF is still the preferred implementation.
 - Independent** - For EGP from campuses in a single routing domain in which, for administrative or technical reasons, the campus is to remain independent of the regional MAN, Border Gateway Protocol v4, (BGP 4) is the preferred solution.

The integrated OSPF implementation will require a detailed implementation plan by the network administrators at the MAN or regional ITSC. Each CAN will be assigned one or more OSPF areas. Connections from the CAN to the MAN IP service provider(s) must be supported by a route redistribution scheme. Within the MAN itself, the ITSC should provide the route redistribution between the separate autonomous networks.

Autonomous networks are considered in this architecture to be stand-alone IP networks and for routing purposes, have their own routing domain. These domains are managed separately and have their own IGP (OSPF). For EGP, autonomous networks will interconnect using BGP 4.

Exchange of routing information between IGP and EGP should only be for IGP within the routing domain of the individual EGPs. EGPs should connect with EGPs in other regions and exchange information about end systems only – they should not exchange OSPF information. Networks announced via BGP 4 should be explicitly configured to ensure high network stability.

C.3.6.4 IP Security Overlay

The DON IP overlay shall use the DON ATM overlay to provide the required confidentiality, integrity, and reliability/robustness. A network intruder can potentially impose an instantaneous and effective denial of service by maliciously reconfiguring routers. The following security mechanisms are included to enable management (and remote management) of routers and to provide protection from spoofing attacks leading to denial of service:

- Exchange of routing table update information between DON IP overlay routers shall use cryptographic authentication mechanisms as specified in the DON ITSG Section 3.4.1.4.
- The IP overlay must provide mechanisms to ensure that CAN network components can only be managed by authorized administrators. At a minimum, network components that are remotely managed must feature non-spoofable authentication mechanisms. It is likely that this management will be accomplished in-band across the IP overlay from a regional ITSC.
- Optional network security mechanisms should be available at the CAN level of the IP overlay to allow for increased security for high value assets or to support command specific security requirements. These optional mechanisms should include:
 - Network Intrusion Filters (NIFs) to provide optional boundary level protections between an organization or CoI and the rest of the DON IP overlay.
 - Network Access Controllers (NACs) to provide a basic level of access control over network connections based on an organization's local security policy.
 - Virtual Private Network (VPN) encryption to allow CoIs to enforce a strict need-to-know separation between their intra-CoI information and the rest of the DON IP overlay. This is normally accomplished by using a COTS VPN encryption mechanism. See the DON ITSG Section 3.4.1.2 for guidance on selection of COTS VPN mechanisms.
- Bringing contractor network connections through the CAN is discouraged. These connections should comply with the guidance in Chapter 3.7.6 and should be negotiated through the ITSC.

C.3.7 Switching and Routing Overlay

Optimized routing across the ATM network will be done using PNNI. Physical routers equipped with ATM interfaces will be necessary for routing between the switched ATM network and existing non-ATM networks. Servers or workstations running native ATM between the desktop and ATM backbone will have multiple connections to the ATM backbone through an ATM workgroup switch.

The CAN connection to the MAN is a switch. To support network security, a firewall router with an ATM interface to the CAN backbone is required at the MAN access point.

The campus infrastructure should be populated with a variety of multi-function hubs that have integrated switch/router functions to enable ATM switching, LANE or MPOA for legacy networking environments, and Switched Ethernet for VLAN support. Interoperability among VLAN vendors is currently the exception rather than the rule, so care must be taken to ensure vendor commitments to embrace available standards.

A combined routing and switching approach will allow work-groups to receive faster network performance with routing invoked only when logical network segmentation is necessary.

C.3.8 Voice Overlay

The integration of voice and data over a single network infrastructure is an important goal of the DON ITI architecture. There are multiple integration levels that should be considered.

For example, a single cable plant can be more flexible to configure and use, can be easier to maintain, and provides a reduced life cycle cost. Category 5 UTP cabling should be used to support both voice and LAN because Category 5 cable is interchangeable (Category 3 telephone cable is not). But this cabling decision alone does not achieve integration – there are many other decisions that must be correctly made to connect computers and hubs to the CAN and telephones and PBXs to the cable.

The frame-based LAN technologies most used in campus LANs today – Ethernet, Token Ring, and FDDI – support voice by digitizing it and placing the bits in IP datagrams. Over a heavily loaded network, the quality of service is less interactive than the network user is accustomed to on circuit switched voice circuits. Alternatively, ATM provides the real-time response and bandwidth-on-demand necessary for telephone quality voice transmission but does not fully support the enhanced voice services and circuit interfaces that are required for full voice functionality. Both approaches are technically feasible, supported in the commercial marketplace, and interoperable via voice-IP gateways.

A remote switching module (RSM) should be co-located with the core (or primary) campus switch. The RSM serves as a central office for the CAN. It typically provides only basic physical voice connectivity. Enhanced services such as voice mail and conference calls are provided by a central voice switch located in the MAN or regional ITSC. The RSM extends the services and reach of the MAN's centralized Private Branch Exchange (PBX) switch to each CAN.

The RSM does not connect directly to the public switched network or other voice networks, but connects to the MAN or region central voice switch, which in turn connects to these networks. ATM provides the transport from the RSM to the region's central voice switch via the CAN primary switch and the MAN/region ITSC ATM switch. This supports maximum flexibility, use of existing infrastructure, centralization of egress points to minimize circuit costs, centralized trouble reporting operations, and easy expansion.

At the campus, the RSM will use load sharing for routing through dual load-sharing ATM devices. In the event of a failure at the switch, the RSM should have stand-alone capabilities until the failure is resolved. Therefore, sufficient redundancy should be built into the CAN architecture to provide the required level of service.

The campus RSMs will be programmed to perform Least Cost Routing functions regardless of the digits dialed by the user to accomplish the most efficient transport egress for the call.

- If a call is destined for a local destination, the switch will use first choice, high usage DoD trunks to the Local Exchange Carrier (LEC) central office for completion.

- If the call is destined for another campus tied to the MAN, the call will be routed to the hub switch for termination to the proper trunk group for the called location, thereby eliminating all LEC toll charges.
- If the call is destined for termination to a facility outside the MAN, a route to the hub switch and an FTS trunk group will be chosen to minimize the cost of this long distance call. Similar routes will be programmed to handle DISN oriented calls from the RSM.

Each CAN voice instrument is directly linked by Unshielded Twisted Pair (UTP) to an RSM port. The cable plant for the premise distribution of voice circuits parallels that of the data circuits. This architecture does not attempt to implement ATM voice or data to the desktop or LAN/enclave wiring closet. However, this capability is supportable within this architecture and it is anticipated that eventually voice and data will coexist on the same campus ATM cable plan.

In the long term, voice and telephony over ATM (VTOA) will provide support for legacy voice services to an ATM terminal. It is analogous to LANE and MPOA in the data environment. As performed on ATM SVCs, VTOA provides improved bandwidth effectiveness (and potential cost savings) because the voice calls are switched directly. At the same time, voice quality is provided through voice adaptation techniques and QoS support.

The voice architecture should take maximum advantage of commercial off-the-shelf technology and existing DoD-sponsored development programs.

C.4 Functional and Performance Specifications

Traditionally, campus networks have been described in terms of design guidance, not in terms of service levels or outcomes. In other words, the government has specified the precise implementation architecture. It is clear that service levels are an integral part of planning and implementation and are necessary to evaluate and determine alternative solutions.

The CAN template will be measured and evaluated under five categories: security, functionality, interoperability, performance, and cost. A level of service and in some cases, the specific technologies required, should be described for each. The following prescribed specifications are provided as a guide; actual requirements may warrant adjustment of these values, but should be based on solid documentation.

C.4.1 Security

A number of security mechanisms are allocated to the CAN. The first layer of defense is to use object level security and the PKI infrastructure as the key exchange method. This places the security focus on the data rather than on the networks through which the data flows and the computers it is stored in. Secure e-mail is one instance of object level security. Cryptographic equipment is required for the protection of classified information.

Information protection must be provided by information security. This architecture details the mechanisms that must be supported.

- Hardened

Level of service: Must provide sufficient level of protection for denial of service and intrusion detection.

Technology: The Campus Network will use a layered information protection system as follows:

1. End users use object level security (such as secure e-mail) as the first line of defense. This provides confidentiality, authenticity, and non-repudiation features to the level of the supporting PKI trust model.
2. Virtual Private Networks provide secure enclaves at varying levels of granularity and these, along with object level security, should be used when available. Because some applications such as FTP are not supported by object level security implementations, this helps to mask some traffic analysis parameters and to provide greater application-level flexibility.
3. Firewalls are located in the DON intranet at the external network interfaces of the ITSC/MAN. Campus networks shall not install bypass routing (such as a router interface to an external network) around the DON firewalls except when they are configured and managed by the ITSC. A campus network may find it necessary to firewall itself from the rest of the Naval intranet by placing a firewall at the CAN/MAN interface.
4. Link level encryption may also be added to links in the CAN when higher levels of protection are required.

- Survivability (specifically relating to security) (see also Performance)

Level of service: Intrusion detection, denial of service, vulnerability to service attacks

Technology: These network control functions are performed at the ITSC NOC, not at the campus level.

C.4.2 Functionality

The capabilities required of the CAN infrastructure that are necessary to effectively and efficiently support the operational mission and requirements must be clearly defined.

- CAN connectivity

Level of service: Each CAN is provided access to two geographically-separated MAN switches.

Technology: N/A

- ATM

Level of service: If applicable, use as required to support mission requirements (voice, video, and data) of the CAN. ATM provides robustness and potential for growth.

Technology: End-to-end SVC

- Switch functionality

Level of service: Must support constant bit rate QoS.

Technology: Non-blocking and provide separate queuing for different QoS classes. Switches use PNNI. Requires NSAP to IP mapping for support of IP.

- IP

Level of service: Ubiquitous end-to-end IP connectivity is supported across the enterprise. It requires full access to the global Internet. Campus routers are connected to the MAN via a full Virtual Circuit mesh.

Technology: N/A

- Router functionality

Level of service: The CAN routing architecture is integrated with the MAN and with the fleet's single routing domain. IP router functionality includes minimum hops between end points and scalability.

Technology: Routers communicate with each other inside the CAN using OSPF and outside the CAN using BGP 4.

- Availability

Level of service: Accessible to all campuses, MANs, and external customers.

Technology: N/A

C.4.3 Interoperability

The components of the network infrastructure must interconnect and efficiently and effectively communicate signaling and other information transfer data.

- Manageability

Level of service: The selection of critical networking components such as switches, multiplexers, and router components of the network must be controlled so that they are remotely manageable by the ITSC to support availability, performance, and cost. This is essential to control cost.

Technology: Components are required to have network management agents to enable remote network management.

- Service Delivery Points

Level of service: Must support the full suite of ITSG-cited ATM protocols and addressing schemas.

Technology: N/A

C.4.4 Performance

The CAN service must be of sufficient quality to meet information transfer requirements to support the Naval mission.

- Bandwidth

Level of service: The bandwidth should be appropriate to the customer's information requirements and be readily scaleable. This is interrelated to availability and survivability.

Technology: Bandwidth is dependent upon technology selected and is discussed elsewhere in this template.

- Delay

Level of service: the cell switching delay on a CAN switch will not be more than 20 us under a load of 80 percent of the interface rate (for example, OC-3c) of both input and output port. The maximum delay variation/switch shall be less than 2 us under the same load condition as above.

Technology: N/A

- Latency

Level of service: the average latency in completing a call setup for a two level peer group must be less than 100 ms/hop.

Technology: N/A

- Network availability. The prescribed level of service should apply to the campus backbone and to critical servers (both network support and user services) that are attached to the campus backbone.

Level of service: It is neither practical nor necessary to provide high availability (higher than 0.99) for each client on the network, because clients tend to be interchangeable (i.e., they view e-mail from multiple desktops). Similarly, it is neither necessary nor practical to provide high availability at the border of the network.

Technology: Topology provides for no switch or physical circuit being a single point of failure.

- Survivability (specifically relating to performance) (See also Security)

Level of service: Vulnerability to forces of nature, acts of humans (e.g., backhoe, loss of power)

Technology: No switch or physical circuit is a single point of failure

- Quality of Service

Level of service: As noted in the MAN template, automatic virtual circuit crossover should be provided by the service provider to provide continuity of service in the event of a switch failure. Ability to support a number of service classes based on the traffic type, each with an associated QoS parameter.

Technology: N/A

- Mean Time to Repair

Level of service: Premium service: less than 2 hours. A CAN planner/manager should evaluate the tradeoff between repair response and redundant equipment. For example, a second router in a location separate from the first router and on a separate Uninterrupted Power Supply may provide adequate redundancy (and additional capacity) to eliminate the need for a 2-hour maintenance response. The maintenance posture can shift to “next working day.” When making this decision, consider the cost of the second router and associated equipment and whether or not it is greater than the on-call maintenance requirement that it may eliminate.

Technology: N/A

C.4.5 Cost

Implementation cost is an important category for judging value and must be done in the context of the level of service. The principal metric for cost control revolves around people costs. The DON will minimize IT support costs at the campus level (especially the 24-hour watch-standing costs and on-call maintenance costs) by centralizing those functions at the ITSCs and eliminating the requirements for on-call maintenance through redundancy and planning for high availability. This centralization initiative assumes that the ITSC will perform effectively and will not increase the IT infrastructure support burden on end users.

C.5 Evaluating CAN Products and Services

The selection of CAN products and services should be made by an organizational team comprised of representative operational and functional experts.

The determination will be made based on a balanced scorecard approach using as a basis the characteristics defined in the Functional and Performance Specifications in Section C.4. Addressing the proposals of each competing product or service should be in accordance with the criteria in Figure C-5.

| Network Characteristic | Raw Score | Weighting (.xx) | Adjusted Score |
|------------------------|-----------|-----------------|----------------|
| Network Security | X | .30 | X |
| Functionality | X | .25 | X |
| Interoperability | X | .20 | X |
| Performance | X | .25 | X |
| Total | — | 1.0 | X |
| Total Merit / Cost | — | — | X/\$YY |

Figure C-5. Product/Service Evaluation

- Four of the five characteristics will be scored based on a numerical rating system (1-10). The supporting subsets of each characteristic will be evaluated and aggregated to determine the value of the characteristic.
- Also, each characteristic will have a weight factor based on its judged importance to the overall region/MAN/CAN mission. The total of these individual weight factors will be 1.0.
- The numerical value for each characteristic is multiplied by its weight factor, and the total of the five adjusted scores equals the relative merit for the source provider.
- Cost will be viewed as a separate variable. The relative merit will be expressed in a ratio of total merit to cost (e.g, X / \$ YY).
- The team should complete an analysis of the alternative providers and present a recommendation to the regional commander/cognizant decision-making body. Included in the presentation will be a recommendation for the administration and funding of the product or service.

Volume I, Appendix D - Table of Contents

| | |
|--|------------|
| D. ITSC Design Template | D-1 |
| D.1 Purpose | D-1 |
| D.2 ITSC Drivers | D-2 |
| D.3 ITSC Customers | D-2 |
| D.4 Elements/Features/Specifications | D-3 |
| D.4.1 ITSC Connectivity to WAN, MAN, and CAN | D-3 |
| D.4.2 ITSC Security Architecture | D-5 |
| D.4.3 ITSC Levels of Connectivity Service | D-6 |
| D.4.4 Description of ITSC Services Provided | D-7 |
| D.5 Hierarchical Service Structure ITSC and ITOC | D-20 |
| D.6 ITSC Infrastructure Physical Attributes | D-21 |
| D.7 Metrics | D-21 |

This page intentionally left blank.

D.ITSC Design Template

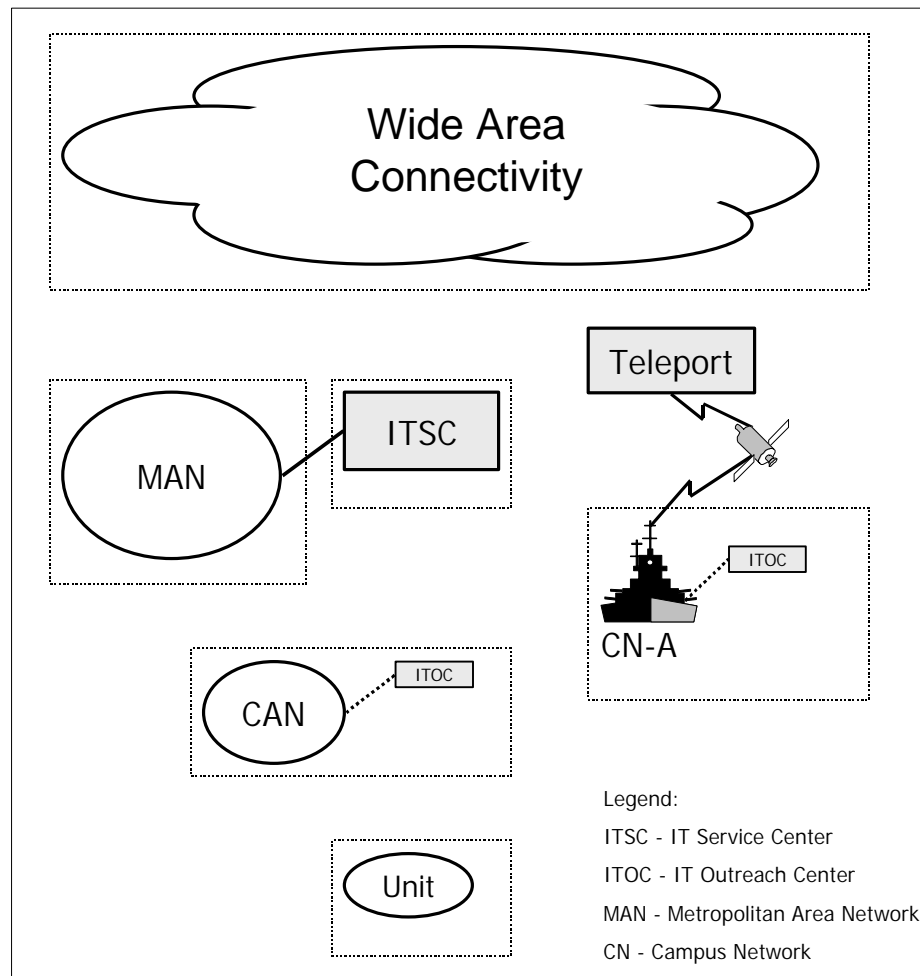


Figure D-1. High Level Components of the TI Architecture

D.1 Purpose

Figure D-1 depicts the relationship of the Information Technology Service Center (ITSC) to the other components of the Information Technology Infrastructure (ITI) Architecture. The ITSCs are operated by the Navy and Marine Corps and provide essential network services.

The collection of ITSCs in the DON forms a cohesive, interconnected, integrated, and interoperable network management community that is dedicated to providing infrastructure support for Navy and Marine organizations. The ITSCs, like the network they support, have the capability to manage voice, video, and data transmission. The ITSC-supported ITI promotes economies and efficiencies by reducing and consolidating redundant services and performing them in such a way that they earn “provider of choice” status among the Navy and Marine Corps user organizations.

In order to support the DON enterprise network, there is a need for cohesive strategy, consistent planning and implementation, and resultant maximum services interoperability and performance.

This architecture must provide sufficient guidance for separate and geographically distant regions to independently implement a Navy- and Marine Corps-determined number of ITSCs that mutually support a DON enterprise network management capability.

This ITSC template establishes a consistent means to address the delivery of these network services, identifies appropriate instances for specific service solutions, and provides a detailed list of implementation considerations. This template is intended for use by ITI planners and implementers and should guide all future network management implementations at all Navy and Marine Corps ITSCs.

D.2 ITSC Drivers

The DON network infrastructure and services must support the diverse and demanding mission requirements of the Navy and Marine Corps. The ITSC architecture solution must satisfactorily meet the following:

- The network infrastructure, and by extension, the ITSC management structure must be supported as a single integrated system consistent with the stated goals of network centric warfare.
- The DON ITSC template must specifically address fleet and mobility issues.
- The location of ITSCs will be determined by criteria that includes the following factors:
 - ♦ geographic distance and distribution
 - ♦ user population density
 - ♦ information technology expertise currently demonstrated
 - ♦ alignment with the operational chain of command
 - ♦ military service mission
- The DON must achieve efficiency through ITSC consolidation but not at the expense of reliability and quality service.
- There must be accountability for ITSC performance in any DON ITSC management solution and customers must be assured of quality service, responsiveness, and best value.
- The ISO model for network management is applicable to organizing the ITSC management functions and will serve as a framework for addressing ITSC-specific responsibilities.

D.3 ITSC Customers

ITSCs must address the network services of all DON customers residing in their areas of responsibility. Generically, the ITSCs must support the following categories of customers:

- Commands not connected via an ITSC-serviced regional MAN with organic servers
- Commands not connected via an ITSC-serviced regional MAN and with no organic servers
- Deployed and underway units
- Commands connected via an ITSC-serviced regional MAN (e.g Tidewater MAN connecting all Campus Area Networks (CANs) in the Tidewater region)
- Information producer commands

- Small units or users on travel
- Interface to the DISN Regional Control Centers (RCCs) and Global Control Center (GCC) under the DII control concept for all DISN access issues within a region

D.4 Elements/Features/Specifications

Figure D-2 provides the outline to address the elements, features, and specifications of the ITSC. The ITSC is described first in terms of its relative position in the WAN, MAN, and CAN topologies and the associated requirements for survivability and security. A concept of service levels provides a means to identify and address the diversity of DON organizational connectivity and service requirements. A description of the services to be provided by the ITSC forms the main body of this template. The relationship of the ITSC to the Information Technology Outreach Center (ITOC) is described. The ITSC physical architecture references the requirements of the ITSC physical infrastructure as provided in the DON Information Technology Standards Guidance (ITSG). Finally, for eventual ITSC implementation, metrics are considered to be essential to gain the necessary acceptance of the major organizations of the Navy and Marine Corps.

| | | |
|---|---|---|
| 1. ITSC Connectivity Relationships to WAN, MAN and CAN | | |
| 2. Concept of ITSC Levels of Connectivity Service | | |
| 3. Description of ITSC Services Provided | | |
| Network Operations <ul style="list-style-type: none"> • Help Desk • Fault Management • Security Management • Performance Management • Accounting Management | Network Administration <ul style="list-style-type: none"> • Capacity Planning • Security Planning • Domain Name System • Dynamic Host Configuration Protocol • Directory • Public Key Infrastructure • Configuration Management | ITSC Supported User Services <ul style="list-style-type: none"> • Network Time Protocol • Email • News/Network News Transfer Protocol • Web • File Transfer Protocol • Remote Access • Multimedia |
| 4. Hierarchical Service Structure of ITSC and ITOC | | |
| 5. ITSC Infrastructure Physical Attributes | | |
| 5. ITSC Metrics | | |

Figure D-2. ITSC Connectivity and Services Outline

D.4.1 ITSC Connectivity to WAN, MAN, and CAN

The elements and features of the ITSC architecture can best be explained in a layered approach. The graphic depicted in Figure D-3 shows the physical connectivity of some of the network components.

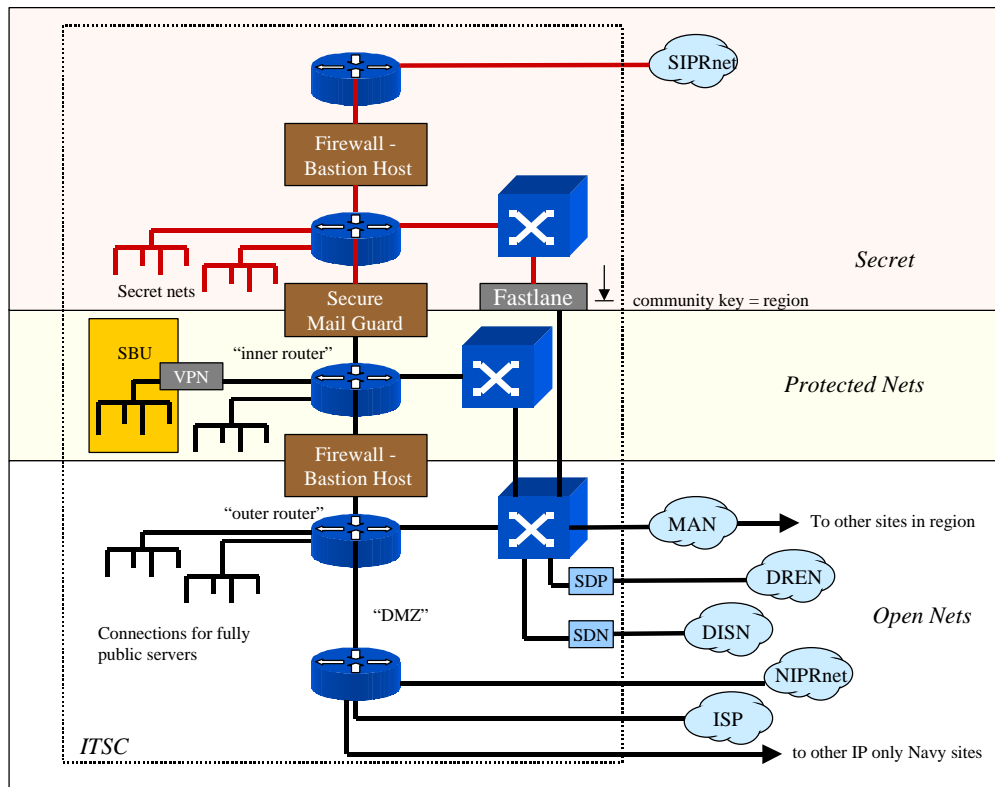


Figure D-3. ITSC Physical Connectivity Layout

The “open networks” shown in the bottom part of Figure D-3 run across the MAN and to the wide area in the clear. They are primarily used for the Demilitarized Zone (DMZ) and peering type functions and for connecting public services.

For unclassified network services, there are two separate logical networks – the Open Nets (outer unprotected connectivity) and the Protected Net (inside the firewall). The campuses will need to access one or the other, or both. The logical separation of these networks is not reflected in this diagram but is accomplished through ATM virtual circuits. For sensitive but unclassified (SBU) network services, VPN technology will be used to encrypt information as it is carried over the networks that are external to the base.

Additional architecture detail can be represented by the logical connectivity that is supported by the ATM Virtual Circuits (VCs) and Virtual Paths (VPs) that ride on top of the depicted physical layout. The VCs and VPs are not shown in the above diagram.

The typical path for a campus to get Internet connectivity is to go through Fastlane encryption. To get SBU network connectivity, the typical path is over the MAN to the ITSC and through Fastlane. In both cases, it is transported through the firewall and onto the NIPRNET and other Internet Service Providers (ISPs) using the IP protocol.

Campuses that have back-door connections to ISPs will not be allowed to connect (via the Fastlanes) inside the firewall. They will connect to the “outer net.”

D.4.2 ITSC Security Architecture

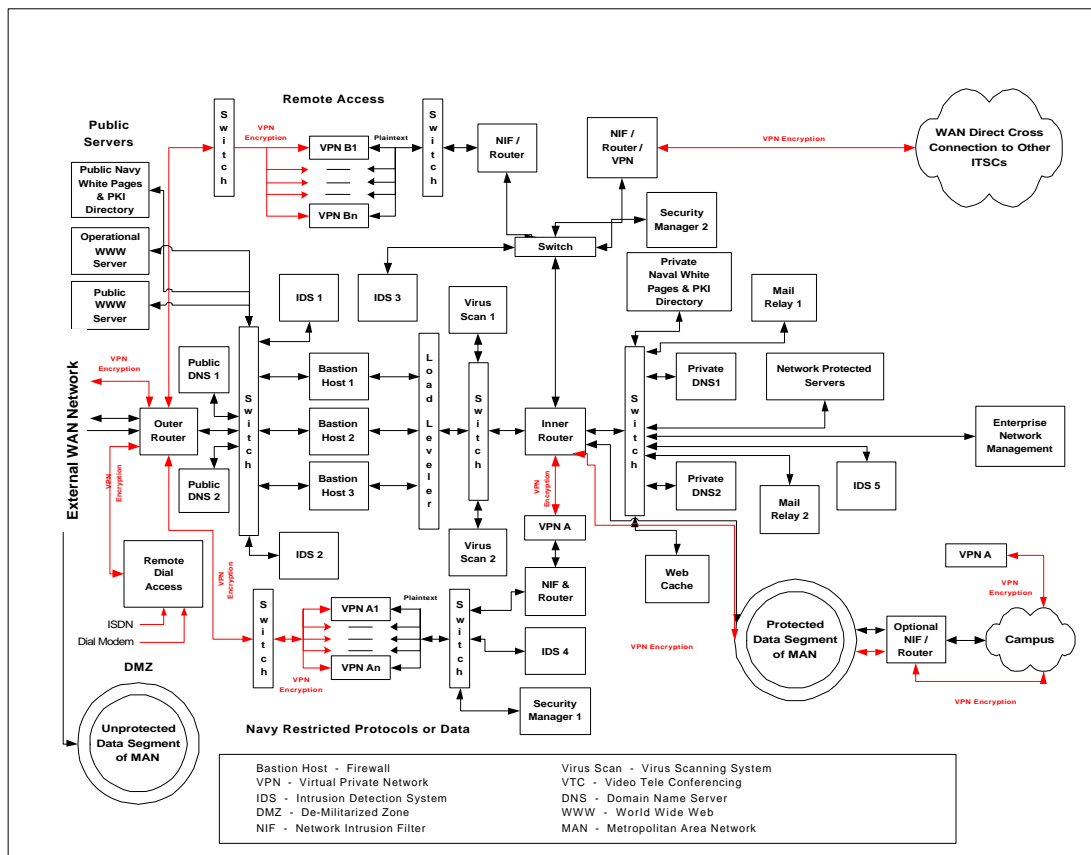


Figure D-4 ITSC Security Template

The template depicted in Figure D-4 describes the various functional components of the ITSC security architecture and how they interconnect. This template should be used for the design of ITSCs.

The template reflects the standard Naval firewall design, which includes the outer router, bastion hosts, inner router, public and private DNS servers, virus scanners, mail relays, and intrusion detection systems. The template also includes an IP “Load Leveler” which performs load balancing across the three bastion host firewalls in order to provide improved scalability.

Most servers are connected behind the firewall in order to protect them from unauthorized external access. The exceptions are the public servers that are connected to the outer router. These servers would be blocked from external access if they were connected inside the firewall. The representative servers shown here are the web servers and a directory (Navy white pages).

For access to the protected services from outside the firewall (remote access), there is a “firewall bypass” using Virtual Private Network (VPN) technology. User who gain access will need to use VPN client software and certificate based authorization. This access from the unprotected side is supported by strong authentication and a fully encrypted path.

For application specific protection where “dangerous” protocols need to pass through the firewall or remote networks need virtual connectivity behind the firewall, there is additional VPN

functionality to meet those requirements. In such cases, the data is broken out into the clear to perform intrusion detection and filtering before re-encryption for distribution in the region.

The MAN is illustrated as having both a “protected data segment” and an “unprotected data segment”. This possibility of multiple segments illustrates that a single physical MAN infrastructure will have a number of separate logical networks to support its multiple requirements. Campus networks will be connected to a MAN on the logical segment that matches its connectivity requirements. In order to connect to the protected data segment, campus networks must completely eliminate all other external connections (back doors). The unprotected data segment is primarily for external connectivity and for peering on the DMZ.

Campuses should still employ intrusion filters and possibly VPN technology even when connected to the protected data segment of the MAN. This added protection supports the DoD “Defense in Depth” approach.

The depicted web cache not only provides added security through its proxy functionality, but also serves a very important function of reducing the consumed bandwidth and the number of connections across the firewall. All web queries to the outside are forced to go through the web cache.

D.4.3 ITSC Levels of Connectivity Service

Levels of service in a MAN or region must address two aspects. First, the level of service for the most demanding customer establishes the highest level of service that must be provided in a MAN or region. To some extent, this level influences the level of service for other users. For example, many applications share the same connectivity, but only a few require the high levels of availability that necessitate 24-hour monitoring and backup connectivity. However, when those services are in place for the demanding users, others may benefit at little or no additional cost. Conversely, premium service requires premium levels of effort and it is obviously not economical to provide that level of service to everyone. The maintenance response for specific applications and associated servers can and should be tuned. It is inappropriate to provide 24-hour application maintenance support for administrative applications that are used only during daytime working hours.

The following generic levels of service and their definitions are established for the DON by this architecture document. Importantly, levels should not be determined by the rank or grade of the commander or official in charge but should be based primarily on the service appropriate to the characteristics of the information required to support the organizational mission.

The following service levels will be consistently applied, as appropriate, for planning and implementation of all connectivity and network services provided by the ITSC.

- Level 1 (High) – Includes operational commands (and certain admiral/general level staffs). Characterized by intense high volume and real-time response network traffic. These commands require 24 x 7 support for all or some network devices.
- Level 2 (Medium) – Includes organizations, staffs, and/or commands that experience interim periods of intense high volume traffic but require a moderate bandwidth and service response. Hours of operation to be supported are normally daytime working hours.
- Level 3 (Low) – Administrative commands (and certain operational and staff commands) that have a moderate to low requirement for network bandwidth and services and require normal response times for performing business-related activities.

D.4.4 Description of ITSC Services Provided

D.4.4.1 Network Operations Center

The Network Operations Center (NOC) services as outlined in Figure D-5 include the network management link to the customer that allows the customer to provide notification of network problems and to receive feedback regarding resolution. These NOC services also include the monitoring services (fault management, security management, and performance management) that enable network managers to monitor the network performance and take corrective actions to improve the network performance in response to both performance degradation and failures. These NOC services support network communications that are mission critical and require timely resolution.

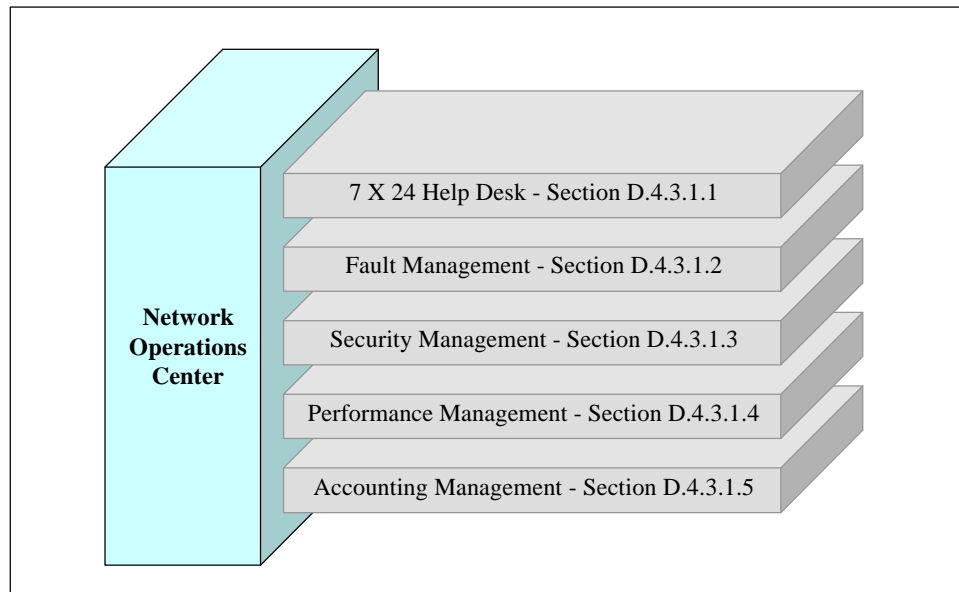


Figure D-5. Network Operations Center Services

D.4.4.1.1 7 X 24 Help Desk

Description: The ITSC Help Desk includes the services associated with end user support. It is responsible for providing multiple resources to solve network computing issues to the client's satisfaction. The Help Desk will be an integrated service provider and will be a single point of contact for all customers who need assistance. By interoperating with all other ITSCs, it will be able to provide tracking, escalation, information sharing, and contingency back up. Specific Help Desk service guidance includes the following:

- Trouble ticket origination and tracking
- Pass off and escalation scheme
- Tracking system database structure
- Recommended fields (schema) for regional trouble tracking databases

Guidance: At a minimum, the ITSC Help Desk services will include the following:

Trouble Ticketing Origination and Tracking.

1. Tracking tools must be integrated with standard fault isolation and reporting tools to enable automatic ticket generation.
2. The Help Desk application must be World Wide Web (WWW)-capable so users can submit and query statuses.

Pass Off and Escalation Scheme

1. There will be a minimum of 3 tiers of support, with tier 1 being the lowest and tier 3 being the highest.
2. All trouble calls are initiated as tier 1 by the ITSC and escalated by the ITSC.
3. Trouble tickets raised to tier 2 and 3 will be fielded on a priority basis with on-call support available outside of normal business hours.
4. Escalation of calls to higher tiers will be performed automatically by mission criticality and case-by-case.
5. Inter-regional pass offs are conducted as follows:
 - ♦ The activity receiving the report will record receipt of the ticket.
 - ♦ The affected region will be notified by the tracking system and must acknowledge receipt.
6. Completed action will be reported to the originating activity.
7. Interaction with cognizant vendors to resolve problems will be conducted by a designated single ITSC entity.

Tracking System Database Structure

1. The Help Desk tracking system will use a single relational database management system.
 - ♦ Initially, each region will have a single RDBMS that is consistent across the enterprise.
 - ♦ Regional systems will migrate to a single enterprise RDBMS.
 - ♦ RDBMS “core” schema must be coordinated across the enterprise.
 - ♦ Core schema must also accommodate regional/local requirements.
 - ♦ This tracking system will also support ITOC requirements.
2. The Help Desk schema will:
 - ♦ standardize the transfer of tickets and retrieval of information among the ITSCs, and
 - ♦ be flexible in its design to allow customization for special local needs without impacting the base application.
3. The system will use a Help Desk software tool to generate and display graphs summarizing relevant performance metrics.
4. Current unresolved issues between the trouble ticketing system and Casualty Reports and Unit Reports should be addressed to minimize duplication of reporting requirements.
5. Base level ITOCs must have access to the tracking system.

Recommended Fields (Schema) for Regional Trouble Tracking Databases

- Help Desk category (what the trouble is)

- Caller identification (who called the trouble in)
- User of the affected system (if other than caller)
- Affected system (software or hardware)
- Affected system component (software or hardware)
- Description of problem (use technical terminology)
- Criticality of work affected (at a minimum, mission critical, mission non-critical)
- Internally-generated tasks (not a trouble call but a task to be done)
 - ♦ Generated adds/changes (e.g., computer move)
 - ♦ Internally-generated tasks from network or configuration management systems
- Location of action (building, ship etc.)
- Location amplification (department, room, deck, etc.)
- History of escalation (what level, when, etc.)
- Special handling (special requirements)
- Action completed (use technical terminology)
- Inventory affected (equipment changed or software modified for forwarding to configuration management)
- Follow-on response warranted (yes or no)

D.4.4.1.2 Fault Management

Description: Fault management includes the management tools that support the availability of user systems by providing the ability to perform fault detection, isolation, and correction. This is the means of promptly notifying the Help Desk of failed equipment, which is an essential component of high availability networking. Fault management includes:

- Monitoring and collection of statistics on traffic conditions and use so potential faults can be forecasted and avoided. This is done by the network manager, who polls various agents for these statistics and assembles the resultant data for viewing (and logging for trend analysis).
- Alarms that warn of threshold conditions on the network that may cause failures. The alarm thresholds are set with SNMP sets and are triggered by SNMP traps.
- Alarms that warn of performance degradation on servers, switches (networks and phone), routers, and area network links. Also, alarms that warn of resource use problems such as low disk space on a server.

Fault management (the first of three ITSC network monitoring functions) allows the NOC to centrally view fault indicators on managed devices remotely and to make corrective adjustments. It requires network management agents in deployed network equipment (e.g. network management agents in routers, hubs, switches, UPSs, as well as end systems and software processes) that the manager interacts with to collect, fuse, filter, and display data. Such interaction requires devices that support enabling standards (e.g., SNMP 2-3 (Simple Network Management Protocol), RMON-II (Remote Monitoring Protocol), MIB-II (Management Information Base).

Guidance: At a minimum, ITSC fault management services will:

- Provide concise and in-depth views of network connections in a graphical format to provide the ability to evaluate network performance, preempt network disruption, and anticipate network growth or realignment.
- Monitor network capacity and use to project future expansion needs. As growth occurs (e.g., new workstations, printers, routers, web servers, and other devices), these tools provide an indication of corresponding required network enhancements.
- Identify network components that exceed set thresholds. Data trends are generated to provide analysis of network condition and developing trends.
- Consolidate network topology into simplified maps to enable generation of multi-level reports that enable robust analysis for optimizing performance at each level. Complex networks are easily viewed to recognize potential bottlenecks and balance network resources for optimal efficiency.
- Forward, by agent, the specific network events from appropriate collection sites to ITSC to ensure that events get appropriate attention. Event notification is by color-coded Graphical User Interface (GUI). Events include:
 - ♦ Equipment failure
 - ♦ Communications loss
 - ♦ Overloaded components
 - ♦ Improperly configured components
- Enable monitoring and notification of local specific events necessary to support specific environments.
- Provide the capability to isolate monitoring of carriers' frame relay, the local telephone companies' circuit, the customers' equipment, and users' applications. One means of accomplishing this is for the local telephone company to export the appropriate SNMP data to the NOC's network manager.
- Enable access to all management information at a single console. All pertinent data will reside on a single ITSC repository.
- Resolve problems and perform actions automatically whether the network is up (in-band) or down (out-of-band).

D.4.4.1.3 Security Management

Description: Security management includes assessment of the managed infrastructure's security posture. It also includes resistance and detection of intrusions and other information protection infractions. Security management will:

- Monitor normal performance of network components and attached end systems. The first indication of a security violation is often abnormal performance.
- Positively control firewall configurations.
- Operate intrusion detection software (most implementations examine packets at firewalls screening for villains).

- Document intrusions in an evidentiary manner suitable for supporting possible law enforcement action.

Security management (the second of three ITSC network monitoring functions) allows the NOC to centrally view security indicators on managed devices remotely and to make corrective adjustments. It requires network management agents in deployed network equipment (e.g., network management agents in routers, hubs, switches, UPSs, as well as end systems and software processes) that the manager interacts with to collect, fuse, filter, and display data. Such interaction requires devices that support enabling standards (e.g., SNMP 2-3, RMON-II, MIB-II).

Guidance: At a minimum, ITSC Security Management services will:

- Monitor existing systems on a 7 X 24 basis to detect insecurities such as spoofing or break-in attempts.
- Monitor network traffic to detect denial of service attacks such as Syn-floods or Smurf attacks.
- Monitor critical processes to ensure that they are not replaced by eavesdropping versions of the processes.
- Maintain intrusion-resistant integrity on network equipment (e.g., proper password control for routers).
- Maintain end system security resistance commensurate with the end system's use. An example is a password sweep to detect easily-breakable passwords.

D.4.4.1.4 Performance Management

Description: Performance management enables informed network performance decisions based on three major categories of performance data: network or site location, physical access circuits, and virtual circuits services. Also included is demand management, which addresses mission and other contingencies that arise that impact normal IT operations and decision-making. For instance, normal IT decisions might be overridden in the event of a resource crisis or other non-IT priority. Specifically, performance management includes:

- Guidance for establishing a plan of monitoring multiple protocol layers and alerts to network operations when performance deviates from normal.
- Corrective actions based on early warning signs to keep network operation at peak efficiency.
- Selective viewing of WAN events by customer network, site, or priority level to allow continuous assessment of performance.

Performance management (the third of three ITSC network monitoring functions) allows the NOC to centrally view performance indicators on managed devices remotely and to make corrective adjustments. It requires network management agents in deployed network equipment (e.g., network management agents in routers, hubs, switches, UPSs, as well as end systems and software processes) that the manager interacts with to collect, fuse, filter, and display data. Such interaction requires devices that support enabling standards (e.g., SNMP 2-3, RMON-II, MIB-II).

Guidance: At a minimum, ITSC performance management services will:

- Report on health of the network on a continuous basis.
- Monitor and tune existing systems on a 7 X 24 basis to ensure optimum use of hardware resources and ensure that agreed performance and throughput levels are maintained. These service levels can be monitored by the ITSC on an exception basis.

- Use performance management to respond to short-term needs and use capacity planning to respond to long-term needs using modeling tools.
- Monitor the network for potential or impending bottlenecks and congestion. Examples of indicators include the following:
 - ♦ Contention-based Ethernet – loading beyond 30 percent
 - ♦ FDDI – loading approaching 96 percent
 - ♦ Telephony – congestion in voice switch resulting in inordinate busy signals
 - ♦ ATM – to be determined
- Report on the server load balancing for servers such as Domain Name System, Dynamic Host Configuration Protocol, web, and Network Time Protocol using diagnostic tools.
- Use performance management data to update capacity planning models on an annual (minimum) or as-needed basis.
- Provide for contingency plans to react appropriately to arising Performance Management information and to allow stabilization of the networks in crisis conditions.

D.4.4.1.5 Accounting Management

Description: Fee-for-service is an important element of network management. Instances in which agencies other than DON obtain network services from organizations within DON must be accounted for in a satisfactory billing arrangement.

Guidance: This will be a topic for future guidance.

D.4.4.2 Network Administration

Network administration, shown in Figure D-6, includes a number of services that support planning for future requirements and administration of network-related functions of a less time-critical nature (e.g., Domain Name Service). These functions are associated with the traditional Network Information Center (NIC).

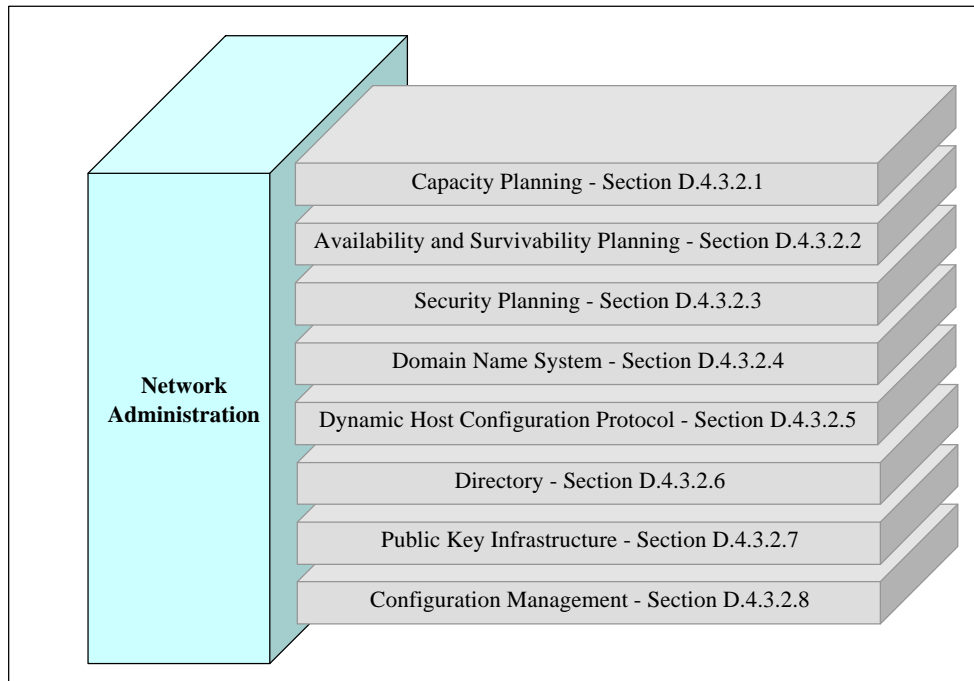


Figure D-6. Network Administration Services

D.4.4.2.1 Capacity Planning

Description: Capacity Planning, which is defined as using modeling to plan changes to the infrastructure, estimates the volume and variety of the workload. It supports analysis of results to ensure that service levels can be maintained and that no major bottlenecks occur.

Guidance: At a minimum, ITSC Capacity Planning services will:

- Use workload forecasts to produce estimates of IT infrastructure requirements to meet service objectives.
- Maintain the following items in a logical information repository:
 - ♦ topology data for the environment,
 - ♦ inventory of equipment,
 - ♦ user workload profiles, and
 - ♦ baseline models representing current resource use.
- Generate forecasting scenarios (simulations) from the repository to plan for meeting service level objectives.
- Use capacity planning methodology to forecast the performance of an active or a proposed system.
- Create a “baseline” of the existing environment that reflects the performance of the infrastructure. The model must be broken into sufficient detail to permit development of resource profiles for each system contained within the model.
- Use capacity planning information to:
 - ♦ estimate future workload growth/change,
 - ♦ forecast the likely use of hardware resources as workload varies,
 - ♦ forecast expected performance and throughput levels for that system, and
 - ♦ project workload requirements and translate them into demands for IT resources.
- Accept information for modeling software from standard ITSC network management tools.

D.4.4.2.2 Availability and Survivability Planning

Description: Availability and Survivability include planning for alternate/redundant connectivity, distribution of routing and switching components, and network monitoring.

Guidance: At a minimum, ITSC availability and survivability planning services will:

- Estimate mission critical systems and plan redundant connectivity (i.e., without single points of failure) to support them.
- Examine existing infrastructure for creeping requirements. Networks tend to attract mission-critical requirements over time that did not exist on the original installation.
- Ensure that installation of networking equipment eliminates or minimizes single points and common cause failures. Examples include:
 - ♦ provide separate UPSs for duplicate equipment (e.g. routers),
 - ♦ ensure that alternate connectivity enters premises through different cable trenches than the primary connectivity, and
 - ♦ places routers in different rooms/buildings so that a fire will not disable both

- Ensure that networking equipment has fault monitoring agents so that failures can be monitored by the NOC. Planners/implementers must ensure that mission critical end systems and applications have SNMP agents for monitoring purposes.
- Plan for appropriate backup power and power distribution for extended electrical outages. For example, when diesel generators are used as a second tier behind uninterruptable power supplies, the generator power must be distributed to routers, hubs, and switches that may be fed by several different power panels.
- Arrange backup peering arrangements between ITSC NOCs so that the failure of an ITSC can be compensated for by neighboring NOCs.
- Plan exercises to test the system.

D.4.4.2.3 Security Planning

Description: Security Planning (or Information Protection) includes planning for redundant and complementary security mechanisms that provide information protection consistent with the mission.

Guidance: At a minimum, ITSC information protection management services will:

- Provide guidance for application developers that incorporate security features in a system early in the planning cycle. Ensure that application developers understand utility-of-object level security that uses the Public Key Infrastructure (PKI) support structure.
- Examine current and planned applications for unintended security leaks. For example, does an application inadvertently make network vulnerabilities visible to an eavesdropper?
- Plan exercises to test information protection ability. Such exercises should test the attrition capability of multiple layers of the defense, not just the ability to withstand attacks on one feature.

D.4.4.2.4 Domain Name System (DNS)

Description: DNS is the service that translates domain names to IP addresses and vice versa. A domain name is a mechanism to give unique names to network devices so that users need not remember their numerical IP addresses. The service is implemented as a hierarchical distributed database and is accessed using a client/server model. The server component of DNS is the subject of this discussion.

Guidance: The DNS services description, as provided in Chapter 4, is fully applicable for ITSC implementation.

D.4.4.2.5 Dynamic Host Configuration Protocol (DHCP)

Description: DHCP supports a number of features that are useful to customers of the network. They include:

- Automatic assignment of IP addresses and configuration data such as netmasks and Domain Name System (DNS) servers.
- Immediate assignment of addresses that are actually used, which results in economic use of IP addresses because unused addresses are reclaimed for reuse.
- Easy system administration because of automation and the ability to administer these details from the server rather than the client.

DHCP is not intended to support mobile users, but it is valuable in supporting dynamic laptop network attachment at remote locations.

Guidance: At a minimum, ITSC DHCP services will encourage use of DHCP to provide some elements of mobility within the campus.

While DHCP is not intended to support mobile users, it is also valuable in supporting dynamic laptop network attachment at remote locations.

D.4.4.2.6 Directory

Description: Directory Services provides a phone book “white pages” function and offers a repository for other information such as phone numbers (office, fax, pager, mobile, and Secure Telephone Unit (STU)), e-mail addresses, and mailing addresses. Additionally, the service is expected to be used for other information about individuals, including passwords, digital certificates, and emergency contact information.

Newer client applications are becoming “directory aware” and use standard protocols to locate information. The dependence of these client applications on directories is increasing its importance in the technology infrastructure.

Guidance: The directory services description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.2.7 Public Key Infrastructure (PKI) Administration.

Description: Public key availability is critical to several security applications, including secure e-mail, secure SNMP management, and secure web service. ITSCs will maintain a public key distributed database that is authentic and complete and will include the public keys of all Naval personnel.

Guidance: At a minimum, ITSC PKI services will support the following.

- Commands will generate public and private keys at their level.
- The private key is given to an individual (e.g., by floppy disk, smart card) or to a command custodian.
- The public key must be transmitted, along with the personal data, to the ITSC’s directory database with an accountability chain that precludes spoofing.
- The ITSCs both maintain the authenticity of the directory database and share/propagate the database among other ITSCs.
- Any user (including those outside the DON) must be able to get the public key of any Naval personnel or command from the database with the confidence that the key is authentic and in accordance with the established trust model.
- Initially, the trust model supported will be similar to that used for military ID cards. Other trust models (such as financial transaction warrants) will be implemented later.

D.4.4.2.8 Configuration Management (CM)

Description: Configuration management will be performed by automated systems throughout the regional ITSCs. The primary purpose is to develop a robust system for managing organizational IT resources. The configuration management system will include inventory information on hardware, software, and associated supporting data. The CM process will support

assessment of implementation alternatives, change management, and interoperability. Use of the CM tool should be intuitive, be supported by fill-in-the-blank screens, and be free from cryptic computer syntax. There are other CM functions that are performed within the IT infrastructure (e.g., security, applications) which are addressed under their respective subject areas.

Guidance: At a minimum, ITSC CM services will:

- Implement a standards-based hardware and software inventory.
 - ♦ Consolidate existing hardware and software inventories in a standard RDBMS.
 - ♦ Provide a system with query capabilities.
- Provide a push/pull distribution system to allow centralized identification and issuance of software upgrades.
- Provide centralized administration of:
 - ♦ minimum configurations for determining and initiating required service and
 - ♦ specific service levels based on individual user requirements, including
 - basic IP,
 - telephony,
 - multimedia, and
 - other IT requirements.
- Set minimum requirements for documentation of the CM data maintained, including baseline, changes made, who made the changes, and why changes were made.
- Use a standard tool set across the enterprise with the following features:
 - ♦ Link management devices and management agents with all network monitoring, reporting, and active fault correcting software to develop a standardized interface for data extraction.
 - ♦ Allow integration into asset and service management solutions and be compatible across ITSCs.
- Support comprehensive change management functions, including change simulation, change history, change detection, and notification. This will be done with an automated interface with the ITSC trouble-tracking software for queries.
- Define templates and defaults for most configurations, which will reduce the time needed to perform an upgrade task and ensure correct implementation.
- Ensure that the following CM procedures are consistent across the enterprise, for example, standard CM database fields (Attachment A).
- Define definitions of administrators' permissions precisely. Policies will be uniformly enforced throughout the enterprise, and unauthorized changes will be automatically flagged.

D.4.4.3 ITSC-supported User Services

This section describes the user services that all functional areas require that must be accessible from the network. These network services, as outlined in Figure D-7, are provided by the family of ITSCs and must have a common planning framework and consistent implementation strategy. Some services must be implemented under an enterprise hierarchical plan. These services are fully described in Chapter 4 and are further described in the template in which additional amplification for ITSC implementation may be necessary.

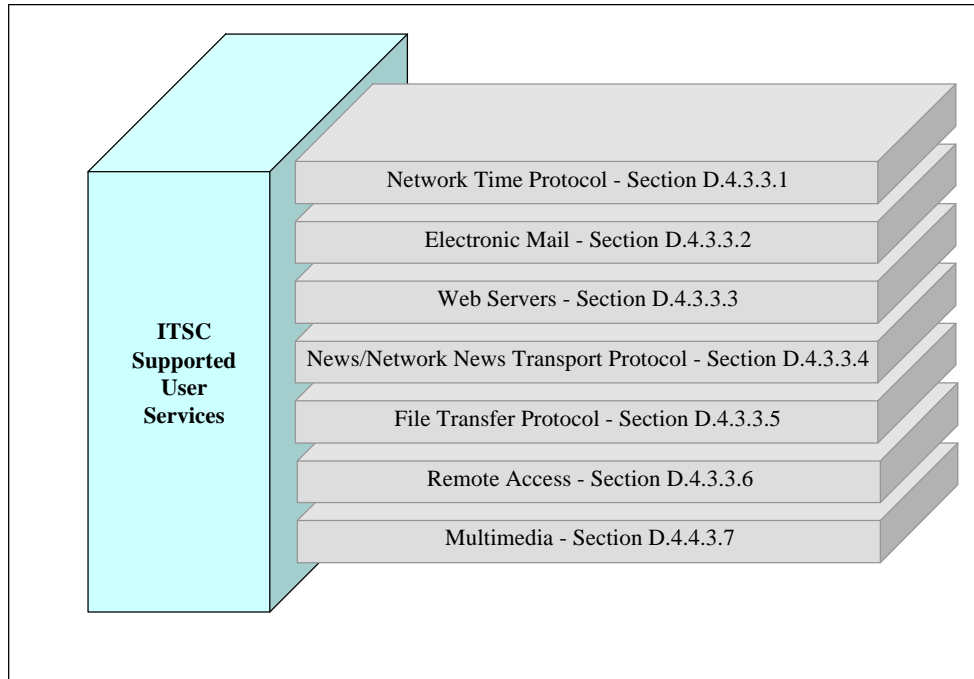


Figure D-7. ITSC Support User Services

D.4.4.3.1 Network Time Protocol (NTP)

Description: NTP services assure accurate local timekeeping so that all servers and services have a consistent time. This is especially important for logging servers, file servers, and some security devices.

Guidance: The NTP service description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.2 Electronic Mail (E-mail)

Description: E-mail is the basic service for interpersonal and organizational messaging which is used throughout the DON. E-mail employs store-and-forward technology that does not provide real-time information exchange, but does provide the capability for high-speed communication of small messages or file transfer. Both individual messaging (E-mail) and organizational messaging (DMS) are supported for the entire Naval enterprise.

Guidance: The E-mail service description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.3 Web Servers

Description: World Wide Web (WWW) service provides one-to-many information sharing throughout the DON and to the public Internet. It is a “pull” technology that allows retrieval (i.e., pull) of shared information from servers. The information is stored in a hypertext transport protocol (HTTP) or eXtended Markup Language (XML) for transport, rendering, and display across an IP network using a WWW browser.

Guidance: The Web Server service description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.4 News/Network News Transport Protocol (NNTP)

Description: NNTP is an information distribution service that provides selective access to “net news.” NNTP varies from an e-mail subscription in that the information content is not stored on “news servers” and is replicated to the degree necessary to provide reasonable local access and performance. Client tools “pull” this content from the news servers and make it available to users for presentation on demand based on the user’s particular selection of “news groups.”

Guidance: The NNTP service description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.5 File Transfer Protocol (FTP)

Description: FTP is used for bulk file upload and download between computers. From a client perspective, a user can connect to a remote computer and either “get” or “put” one or more files as well as perform other simple file manipulation commands. FTP provides a solution for many of the e-mail shortcomings. Large files can be distributed by placing them onto an ftp server and then announcing a pointer to that location so recipients can download needed files at their convenience.

Guidance: The FTP service description as provided in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.6 Remote Access

Description: Remote access provides a modem pool for telecommuters, travelers, and other users to dial up and gain access to the DON enterprise. Regional ITSCs cooperate to publish local access numbers for all metropolitan areas in all regions. This permits frequent travelers to dial local numbers for access to the enterprise. Secure remote access through the Internet is also provided.

Guidance: The Remote Access service description in Chapter 4 is fully applicable for ITSC implementation.

D.4.4.3.7 Multimedia

Description: Multimedia services include video teleconferencing (VTC), video applications sharing, video teletraining, and video and image/graphics file servers. Also covered are VTC application enhanced data services that allow users to share applications/documents and to participate in collaborative activities including video applications sharing, video document sharing, and “white boarding.” VTC allows geographically-dispersed personnel and activities to conduct face-to-face meetings in real time through the transmission of images and sound.

Guidance: The Multimedia service description in Chapter 4 is fully applicable for ITSC implementation.

D.5 Hierarchical Service Structure ITSC and ITOC

The services described in this template are intended for execution by the DON ITSC management structure. Determining the way that these services are distributed is not the purpose of this initial DON ITI architecture document, rather, the purpose is to establish ITSC services and a description of those services. Having defined the functions to be performed, it is then appropriate to determine the structure that can best provide those services efficiently, effectively, competitively, and with accountability to the customers. This ITSC organization is a pivotal element in obtaining a world-class DON enterprise network and appropriate, cross-functional input and attention to this decision is expected.

It is obvious that there will be ITSCs in each of the MANs or regions and these will have extension services at the campuses. A description of these extensions, called Information Technology Outreach Centers (ITOCs), is provided. The ITOC functions are not a duplication of the consolidated ITSC functions. For example, an ashore ITOC may be able to view a NOC's network management map on demand, but it cannot generate its own. Importantly, both the command relationship and the required functionality of the ITOC differ depending on whether the ITOC is located ashore or if it is deployed/afloat.

D.5.1.1.1 Ashore ITOCs

Ashore, base Information Technology Outreach Centers (ITOCs) have electronic access to ITSC operations, administration, and service information and have appropriate tools to monitor, diagnose, and correct network problems as well as computer and peripheral hardware problems. Because an ITOC can be expected to serve several commands at a campus, and because of the need for ITSC/ITOC integrity, it is logical that the ITOC be part of the ITSC command

In the case of ashore units, stable connectivity to the ITSC can be assumed, so the ITOC requirement to directly run NOC operations can generally be dispensed in favor of getting the NOC's management picture when needed. This elimination of local 24-hour watch requirements represents a considerable economy. Similarly, end system administration, especially of servers, can be centralized at the ITSC (even if a server is physically located remotely from the ITSC). By contrast, the Help Desk and applications planning functions (both day-time work functions) can be expected to be larger than on deployed units. The ITOC should present an effective Help Desk "human face" to the end user -- the ITOC contact should be able to work seamlessly with the ITSC Help Desk staff to escalate and resolve problems and provide customer feedback.

D.5.1.1.2 Deployed/Afloat ITOCs

Operational units require unit integrity and thus, a significant degree of autonomy and self-sufficiency. Therefore, the ITOC will be part of ship's company. The ITOC capability is a subset of the ITSCs function and generally includes:

- Network management. This is a subset of the NOC capability. In the case of operational units, they operate their own network management system and monitor agents within the command by forwarding a consolidated picture to the ITSC (this is constrained by bandwidth use over radio-WAN). This function generally requires a 24-hour watch standing network management capability while deployed. When a ship is tied up in port and has established shore-tie network connectivity, it can transfer its NOC watch-standing requirement to the ITSC NOC and secure its own NOC operation.

- End system administration. Both user clients and some servers will be part of a deployed unit's IT organization and will require maintenance in terms of operating system updates, configuration, application software installation, and license management and control. End system administrators will be expected to generate PKI certificates and authentically forward the public key personas to the ITSCs (refer to Chapter 4, PKI, and the directory).
- IT hardware installation, maintenance and troubleshooting. An extreme example is the Marine Corps requirement to set up the IT shop when phasing ashore in expeditionary operations. But even in these relatively more stable installations, additions, changes, and casualty maintenance are required.

D.6 ITSC Infrastructure Physical Attributes

The resultant characteristics of ITSC infrastructure are equally important for the consistency and quality of IT functions and services provided to the Naval enterprise network. The work of world-class service centers must be considered and incorporated to ensure that ITSC support to the Navy and Marine Corps is consistent with mission support expectations.

The standards for power, cooling, security, fire suppression, cable plant, and other related physical characteristics are extensively addressed in the DON ITSG in Chapter 4.

The implementation of ITSCs should not be left to individual discretion or interpretation or to available expertise. One ITSC information source is the National Association of Network Operating Group (NANOG). Information from the NANOG is available at <http://www-personal.umich.edu/~wbn/DataCenterNeedsNotes.htm>.

D.7 Metrics

Selected metrics for the ITSC are outlined in this document and the DON ITSG. Metrics and ITSC are viewed as critical success factors in any DON enterprise solution. Specific feedback from the largest Naval organizations showed a widely-held skepticism that any existing Naval organization could provide an acceptable ITSC service on a world-class basis. For this reason, these organizations are reluctant to align with any enterprise ITSC initiative. Metrics are therefore an important element of describing what specifically is to be provided and a means to ensure that the promised service is delivered.

The cost and performance of the ITSC implementation must be supported by robust metrics that support the ITSG as a best value. The functions of the ITSCs, when established, will be reviewed for alternative sourcing on a regular basis. When performance and cost metrics do not support it, alternatives to performing a function in the ITSC will be evaluated. This is concluded to be an important strategy to avoid substandard service.

Alignment of ITSC services with the needs of all ITSC customers is a primary focus. Metrics should be chosen to provide clear indication of the ability of the ITSC to meet customer requirements instead of the ability to meet internal ITSC priorities.

Industry best practices provide an excellent source for modeling and aligning network services. ITSG Chapter 10.4.6 provides a strategy for selection of the metrics.

This page intentionally left blank.

Volume I, Appendix E – Table of Contents

| | |
|---|------------|
| E. Network Security..... | E-1 |
| E.1 Purpose | E-1 |
| E.2 DON Information Protection Requirements | E-2 |
| E.2.1 Requirements Areas..... | E-2 |
| E.2.2 Protection Mechanisms..... | E-3 |
| E.2.3 Protection Dimensions..... | E-3 |
| E.3 Defense in Depth Approach..... | E-3 |
| E.4 Mechanisms | E-5 |
| E.4.1 Zone 4 Mechanisms..... | E-5 |
| E.4.2 Zone 3 Mechanisms..... | E-8 |
| E.4.3 Zone 2 Mechanisms..... | E-9 |
| E.4.4 Zone 1 Mechanisms..... | E-9 |
| E.5 Infrastructure Security Components..... | E-10 |
| E.5.1 Infrastructure Mechanism – Secure Domain Name System | E-11 |
| E.5.2 Infrastructure Mechanism – Public Key Infrastructure | E-11 |
| E.5.3 Infrastructure Mechanism – Network Management Security | E-11 |
| E.5.4 Infrastructure Mechanisms – Network Security Configuration Management | E-11 |
| E.5.5 Infrastructure Mechanisms – ATM-level Security..... | E-11 |
| E.5.6 Infrastructure Mechanisms – IP Security..... | E-12 |

This page intentionally left blank.

E. Network Security

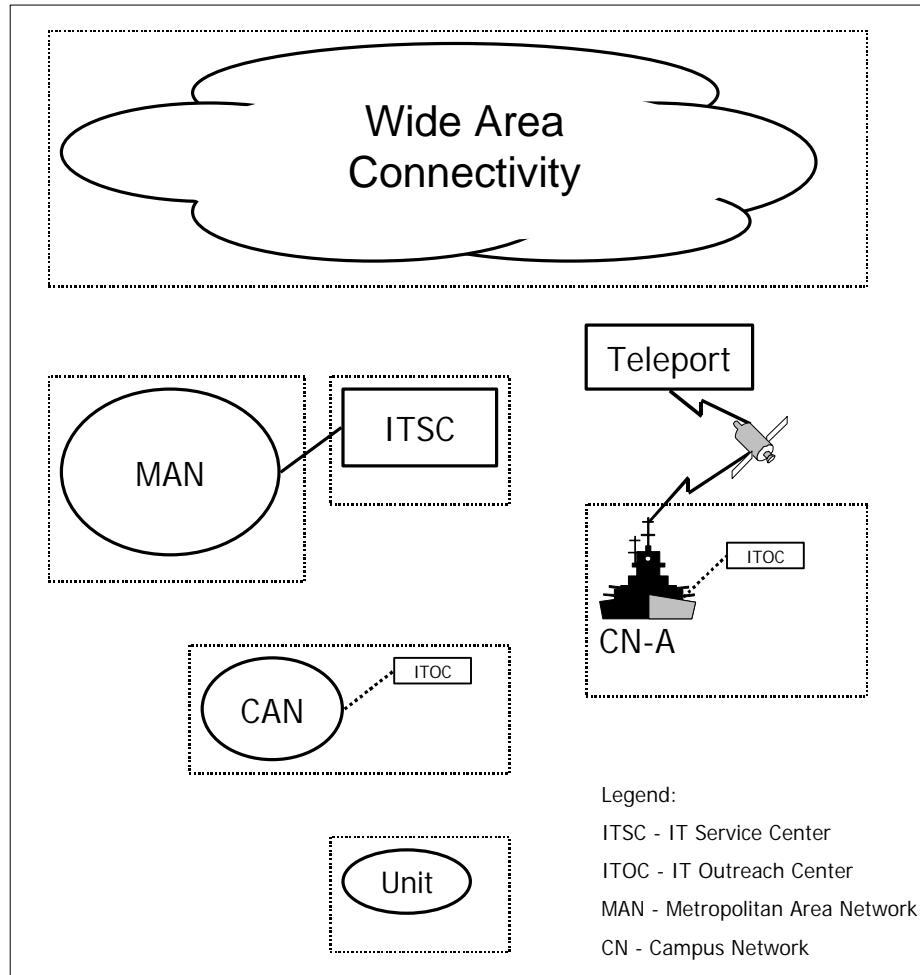


Figure E-1. High Level Components of ITI Architecture

E.1 Purpose

The components of the ITI architecture shown in Figure E-1 were addressed in their respective templates in the preceding appendices. Each template included specific security mechanisms that were appropriate to that level of the ITI infrastructure. Security was integrated into each template to ensure a clear definition of the placement and configuration of the security mechanisms. In this final appendix, security is discussed in aggregate, describing how all of these security mechanisms combine into a “Defense in Depth” security system. ITI Security is addressed in the following three segments:

- The information protection requirements that must be addressed in the Navy and Marine Corps
- The significant infrastructure mechanisms that are used to provide protection at specific points in the technology infrastructure

- A DON defense in depth strategy used to frame the Navy and Marine Corps information protection strategy

E.2 DON Information Protection Requirements

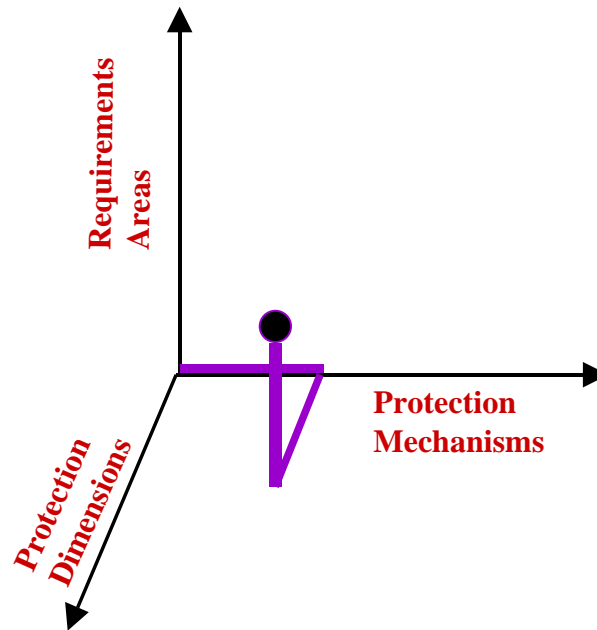


Figure E-2. Information Protection Summary

Figure E-2 depicts the interrelationship of the requirements, mechanisms, and dimensions in a three-axis model. There are five protection requirements areas, six protection mechanisms, and three protection dimensions. The simplified model shown in Figure E-2 masks the innumerable intersections of the three axes that produce nearly 100 instances of specific information protection requirements. The information protection requirements must be defined for each of the requirements intersections.

E.2.1 Requirements Areas

As DON IT regionalization efforts proceed, system designers must ensure that the resulting systems meet certain information protection requirements. The generic information protection classes of requirements for DON information systems are:

- **Confidentiality** – the protection of classified and sensitive but unclassified (SBU) information from unauthorized disclosure
- **Integrity** – the protection of information and information system resources from unauthorized, undetected modification
- **Availability** – the assurance that authorized users will have reliable and timely access to required resources (including information, system services, communication services, etc.)
- **Authenticity** – the ability to determine if information was created or modified by an authorized entity

- **Non-repudiation** – the ability to provide non-forgable proof of a data originator’s identity and non-forgable proof of data receipt

E.2.2 Protection Mechanisms

The following broad categories of information protection mechanisms satisfy the information protection requirements:

- **Encryption** – convert understandable information into unintelligible data for storage and transport in harmful environments and restore this information (decryption) to authorized users
- **Access Control** – control access to system data and resources based on a user’s identity or operational role
- **User Identification and Authentication (I&A)** – securely determine a user’s identity or operational role
- **Malicious Content Detection** – examine incoming data to detect and block malicious content (e.g., viruses)
- **Audit** – record security-relevant events in a protected form (for use in non-real time event reconstruction and in real time intrusion detection)
- **Physical and Environmental Controls** – physically protect and provide for continuity of operations for system components relating to policies, procedures, and mechanisms

E.2.3 Protection Dimensions

To achieve information protection over the DON enterprise, the information architecture is categorized in dimensions that must be protected. The top-level dimensions are listed from more general to specific:

- **Information System** – the actual infrastructure that must be protected against unauthorized intrusion and denial of service
- **Information Domain** – communities of interest (CoIs) within the infrastructure must be afforded freedom to move and process information within a virtual enclave that provides protection
- **Information Content** – information packages themselves must be protected against unauthorized access by untrusted users both while in transit and in storage

E.3 Defense in Depth Approach

Defense in depth is the preferred approach for information protection for the Naval enterprise as regionalization efforts proceed. In defense in depth, information protection mechanisms are applied in multiple, complementary, and redundant locations in a system architecture. The system satisfies its information protection policy while maximizing resistance to attack and minimizing the potential that a flaw in a single security mechanism will lead to a security compromise. Additionally, defense in depth allows for varying levels of permeability in the information protection architecture. For example, more strict security restrictions can be placed on the flow of information from the Internet to a Naval regional MAN than the restrictions placed on the flow of information between two Naval regional MANs.

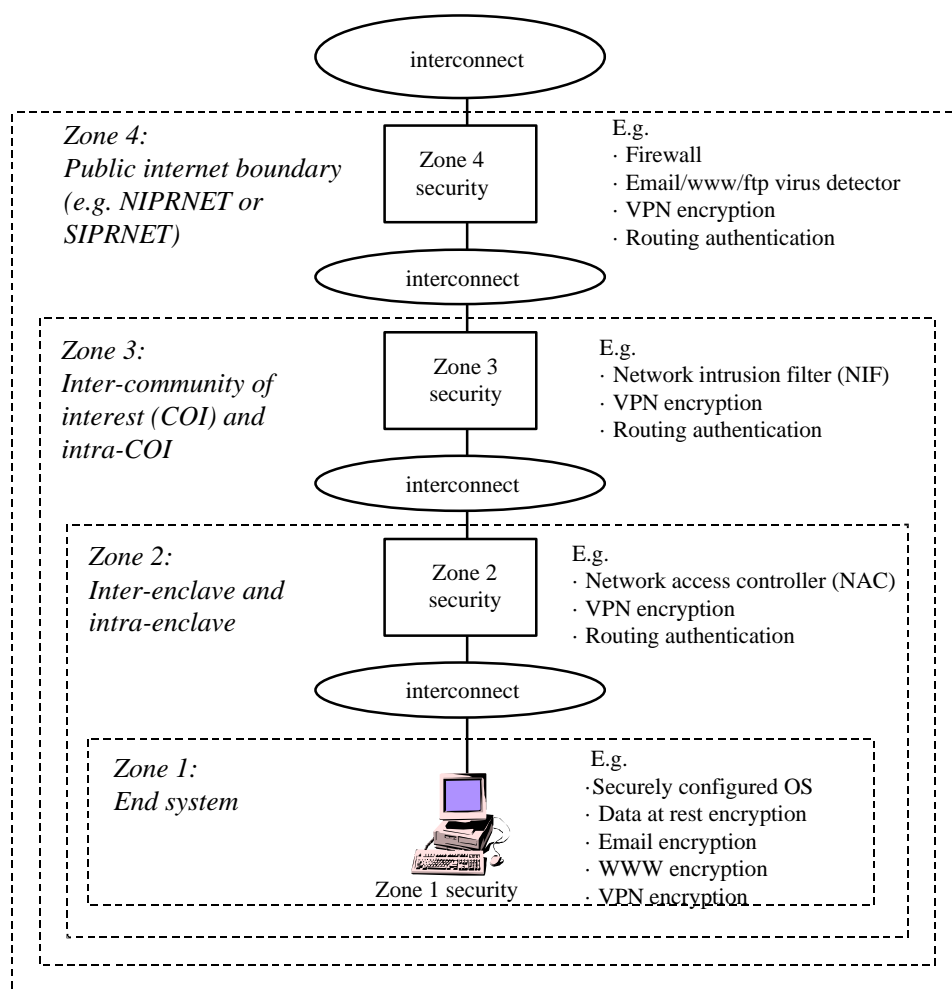


Figure E-3. Generic Framework for Defense in Depth

The defense in depth information protection concept shown in Figure E-3 is directly analogous to naval sea control concepts. Fleet air defense is a representative example. The outer zone is defended by intercept fighters such as F-14s and controlled by E-2Cs. A second layer of defense is the missile zone defended by Aegis cruisers which intercept attackers that are not defeated in the outer layer. Inside the missile zone are the point defense zones where the defensive weapons include chaff, close-in warfare systems, and tactical electronic warfare machinery. Because of the additive effectiveness of these layers, the number of “leakers” that penetrate to the inner zone is less than the capacity of the point defense weapons.

In the Figure E-3 framework, four security layers are defined in the generic framework for defense in depth. The zones of defense may be logical and not necessarily physically separate. Each of the four layers includes appropriate information protection mechanisms selected from the six categories of protection mechanisms and which are required to build secure DON information systems. The most critical of these components is PKI required to support identification and authentication mechanisms and encryption mechanisms that can be applied in the four separate zones.

Figure E-4 summarizes how information protection mechanisms (lower level of granularity derived from section E.2.2) are applied to each architecture security zone and information dimension. A collection of these mechanisms is used to establish the protection needed at each zone. A more in-depth discussion of each information protection mechanism (including appropriate standards guidance) can be found in Chapter 3 of the DON ITSG.

| Mechanisms | Info System | Info Domain | Info Content | ZONE | | | |
|--------------------------------|-------------|-------------------------------------|-------------------------------------|-------------------------------------|---|---|---|
| | | | | 1 | 2 | 3 | 4 |
| Bulk, Link, Source Encryption | C,I,a | C,I,a | C,I,a | ✓ | ✓ | ✓ | ✓ |
| VPN Encryption | C,I,a | C,I,a | C,I,a | ✓ | | ✓ | ✓ |
| Data at Rest Encryption | | C,I,a | C,I,a | ✓ | | | |
| WWW Encryption | | C,I,a | C,I,a | ✓ | | | |
| Email Encryption | | | C,I,a,N | ✓ | | | |
| Digital Signatures | | | C,I,a,N | ✓ | | | |
| Routing Table Authentication | A,a | A,a | | | ✓ | ✓ | ✓ |
| Public Key Infrastructure | | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | | | |
| Network Firewall | C,I,A,a | C,I,A,a | | | ⊙ | ⊙ | ✓ |
| Network Intrusion Filter | C,I,A,a | C,I,A,a | | | ⊙ | ✓ | ⊙ |
| Network Access Controller | C,I,A,a | C,I,A,a | | | ✓ | | |
| Content Security Checking | | I,A | I,A | ✓ | ⊙ | ⊙ | ✓ |
| Operating System Security | C,I,A,a | C,I,A,a | C,I,A,a | ✓ | | | |
| Operating System Configuration | C,I,A,a | C,I,A,a | C,I,A,a | ✓ | | | |
| Password Service | a | a | a | ✓ | | | ✓ |
| | | | | | | | |

C confidentiality a authenticity ✓ applicable ⊙ optional
 I integrity N non-repudiation ☒ infrastructure component
 A availability

Figure E-4. Application of Information Protection Mechanisms to Zones

E.4 Mechanisms

A collection of security mechanisms is used to establish the protection needed at each zone for each information protection dimension. These are applied in the four zones that comprise the Defense in Depth information security strategy for DON.

E.4.1 Zone 4 Mechanisms

Zone 4 information protection mechanisms are employed at the boundary between the DON enterprise and a public internetwork (e.g., NIPRNET and SIPRNET). These mechanisms are most likely located in an ITSC or supporting facility such as a firewall facility (FWF).

E.4.1.1 Zone 4 Mechanism - Network Firewall

The primary Zone 4 information protection mechanism is the network firewall. The network firewall forms the boundary between the networks under positive Naval control (e.g., regional MANs) and networks not under Naval control (e.g., the NIPRNET/Internet and SIPRNET). This firewall selectively allows external entities to pass information to systems inside the Naval

enterprise while blocking many potential attacks (including content security checking or detection of viruses in incoming e-mail attachments and file downloads).

Unprotected modem access defeats the purpose of the firewall and shall be disallowed inside this zone.

Network firewalls should be installed at all gateways to the Internet, NIPRNET, and SIPRNET. These gateways will be located at FWFs, which can be located in ITSCs and/or at other locations within a region for redundancy.

A typical network firewall is the Naval Firewall Security System (NFSS) depicted in Figure E-5. NFSS provides redundant bastion host firewalls, virus checking servers, and a Web cache.

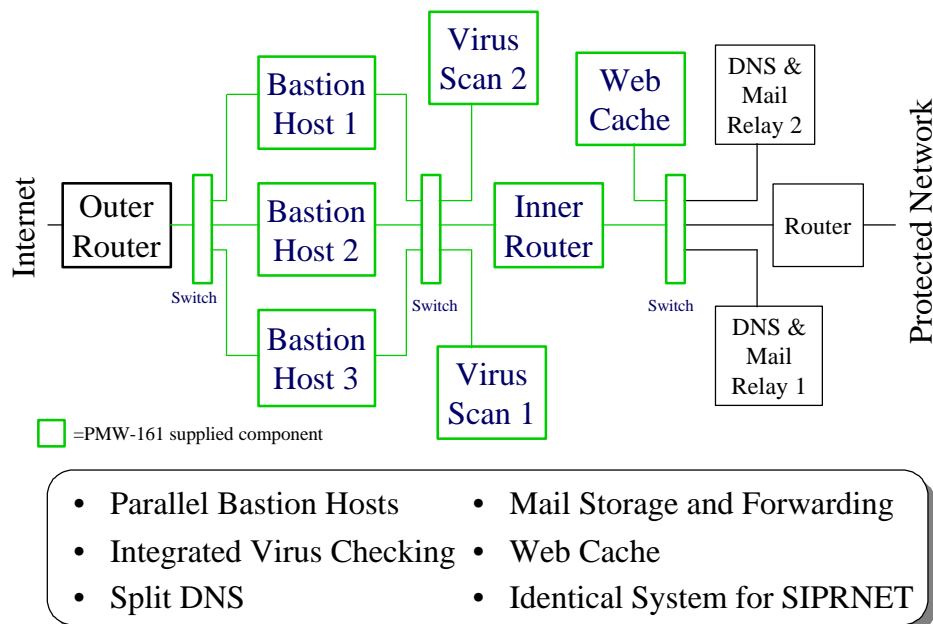


Figure E-5. Naval Firewall Security System

E.4.1.2 Zone 4 Mechanism – Content Security Checking

E-mail attachments, downloaded files, and comparable items should be automatically checked for viruses as they enter the Naval enterprise from the Internet, NIPRNET, and SIPRNET. This check should be performed as part of the network firewall function.

E.4.1.3 Zone 4 Mechanism – Virtual Private Network (VPN) Encryption

VPN encryption can be used to provide confidentiality and integrity for data transmitted across a public internetwork. Additionally, VPN encryption can provide authentication of the remote system that encrypted the data. When integrated into a suitable system architecture, VPN encryption allows secure “tunnels” to be established across a non-secure internetwork. This allows a private intranet to run securely over a public Internet.

When Naval customers can only connect to ITSCs via the Internet, NIPRNET, or SIPRNET, COTS VPN encryption should be used. This encryption should use the standards set forth in section 3 of the ITSG. This VPN encryption should be integrated with Zone 4 network firewalls.

In addition, connections to classified external networks (e.g., the SIPRNET) must be VPN-encrypted if they are not encrypted using NSA-approved bulk encryptors. This VPN encryption must be performed using an NSA-approved in-line network encryptor (INE). The selection of this device must be coordinated with the owner of the destination network (e.g., DISA) and with CNO N643 for requirement validation and central procurement.

E.4.1.4 Zone 4 Mechanism – User Authentication

Under certain circumstances, non-DON users may require access to internal information services (e.g., an FTP server). Additionally, Naval users may sometimes need to connect back to internal information services from external networks (e.g., while on travel and connecting from a contractor LAN). At a minimum, a one-time encrypted password scheme should be implemented as part of the network firewall to authenticate users requesting such access. As an alternative, SSL with X.509 client certificates may be used to authenticate users to either the network firewall or directly to a secure web server hosted inside the network firewall.

E.4.1.5 Zone 4 Mechanism – Intrusion Detection System

An intrusion detection system (IDS) can detect a wide variety of known network attacks by monitoring network traffic and looking for signatures of known attacks. The Fleet Information Warfare Center (FIWC) has the capability to centrally monitor reports from a commercial IDS known as NetRanger. It is expected that NetRanger will be used to establish a DON centrally-monitored intrusion detection capability. NetRanger IDS sensors should be installed together with network firewalls to detect intrusion incoming from the Internet, NIPRNET, and SIPRNET. These sensors should be installed outside the network firewall to enable FIWC to detect potential attacks that are blocked by the network firewall.

E.4.1.6 Zone 4 Mechanism – Bulk (line) Encryption

Connections to classified external networks (e.g., the SIPRNET) must be bulk-encrypted if they are not already VPN-encrypted using an NSA-approved INE. An NSA-approved bulk encryption device must be used. The selection of this device must be coordinated with the owner of the destination network (e.g., DISA) and with CNO N643 for requirement validation and central procurement.

E.4.1.7 Zone 4 Mechanism – Routing Table Authentication

Routing information exchanged with NIPRNET and SIPRNET core routers (under DISA control) should be authenticated with keyed MD5 where possible. This reduces the possibility that an IP route spoofing attack could be used to disrupt NIPRNET or SIPRNET service. The implementation of keyed MD5 routing protocol authentication for NIPRNET and SIPRNET connections must be coordinated with DISA.

E.4.1.8 Zone 4 Mechanism – ATM Address Filtering

Address filtering is another security technology at the boundary between the DON enterprise ATM network and a public or shared WAN network (DISN ATM WAN, Internet, NIPRNET, etc.). Address filtering is used to allow connections to be made only to and from particular trusted addresses and thus precludes general connectivity. However, in the absence of a PKI, it is not clear that address spoofing can be prevented.

E.4.1.9 Zone 4 Mechanism – ATM Cell Payload Encryptors

ATM cell payload encryptors are becoming available, but interoperability between vendor offerings is not likely.

E.4.1.10 Zone 4 Mechanism – Contractor Network Connections

Extending the Naval intranet (WAN, MAN, CAN, or LAN) to a contractor site may introduce new opportunities for security compromise because the contractor site is outside the Naval physical security perimeter. If the connection to the Naval infrastructure is established outside the security perimeter, then no additional security is required. If, however, the requirement is to penetrate the security infrastructure, then it must be done in a manner that does not weaken or compromise the overall security architecture.

When non-Naval networks connect to the Naval intranet inside the firewall, the following issues must be addressed:

- Does the contractor site connect to other networks? If so, these will constitute back-door connections and the connection should not be allowed.
- Can the contractor network be segmented so that one segment is a separate network without back-door connections? If so, this provides adequate protection and enables satisfactory connection to the Naval infrastructure. This segmentation may not be a desirable option for the contractor site.
- Does the contractor site have adequate controls of physical access? For example, are there locked doors? Who is allowed access to the protected network? What are the policies and procedures for establishing a connection to the protected network? Who can use the workstations that are connected to the physical network?

The Naval intranet will provide external access using VPN technology. This allows access to the Naval intranet from anywhere on the global Internet by establishing a secure tunnel through the Internet to the Naval intranet. Access can be controlled by use of an identity certificate provided by the Naval/DoD Public Key Infrastructure (PKI) solution (described in the next chapter). Using PKI, access to the Naval intranet is controlled on an individual basis instead of at the network or device level. Access to the Naval intranet is granted only to those individuals who have a recognized need for network access. Authorized users must still be authenticated to gain access to systems and applications on the network just like any other Naval user.

This VPN solution provides a general solution for contractor access, as well as for anyone (including Naval personnel who travel and/or telecommute) who requires access from outside of the Naval intranet. In this manner, direct connections can be established from outside the security perimeter while preserving the integrity of the security architecture.

E.4.2 Zone 3 Mechanisms

Zone 3 information protection mechanisms are employed to provide optional security protections to high value information, organizations, or CoIs. These mechanisms may be installed at the ITSC or the campus ITOC. Although DON enterprise network resources are protected with Zone 3 mechanisms, hosts and network devices controlled by individual commands and/or CoI must implement their own Zone 3 protection mechanisms, but do not establish security policy for organizations/CoIs.

E.4.2.1 Zone 3 Mechanism – Network Intrusion Filter

The Network Intrusion Filter (NIF) provides an optional level of boundary level protections between an organization or CoI and the rest of a region. This allows for high value assets to be afforded additional protections and/or for CoIs to implement more restrictive security policies than the rest of a region. An NIF is essential when connecting fleet teleports to a region if a network firewall is not in place at the teleport. NIFs may be constructed using bastion host firewalls, stateful monitoring firewalls, tightly configured filtering routers, or intrusion detection systems with active shunning capabilities. NIFs may be used to implement VPN encryption to allow a CoI to be distributed across a region or the enterprise.

E.4.2.2 Zone 3 Mechanism – VPN Encryption

VPN encryption is normally implemented using a NIF if required at Zone 3.

E.4.2.3 Zone 3 Mechanism – IDS

IDS may be optionally applied at the boundary between an organization or CoI and the rest of a region. This IDS should be centrally monitored (see Zone 4). Centrally-monitored intrusion detection is considered essential for fleet teleports.

E.4.3 Zone 2 Mechanisms

Zone 2 information protection mechanisms provide boundary protections between a campus LAN and the MAN or region as well as protections on the campus LAN. These protections are normally installed on LANs.

E.4.3.1 Zone 2 Mechanism – Network Access Controller

A network access controller provides a basic level of access control over network connections based on a site/enclave's local security policy. These controls could include restrictions on incoming connections as well as connections between LAN segments internal to the site/enclave. These restrictions could be based on the source and destination addresses of the IP packet as well as the service type (e.g., SMTP e-mail, telnet, HTTP). A NAC should be implemented using the organic filtering IP routers used to connect the site/enclave to the external world. For ATM systems featuring "cut-through" routing, ATM switches should use NSAP filters to ensure all incoming packets are routed to the NAC for security filtering.

E.4.3.2 Zone 2 Mechanism – VPN Encryption

VPN encryption is normally implemented using a NAC if required at Zone 2.

E.4.3.3 Zone 2 Mechanism – Routing Table Authentication

IP routers internal to the DON enterprise should be configured to use keyed MD5 authentication for routing updates. See Zone 4.

E.4.4 Zone 1 Mechanisms

Zone 1 information protection mechanisms provide the inner-most layer of defense for information systems. The protections are implemented on the actual end systems, including NT workstations, NT servers, UNIX servers, and mainframes. Zone 1 information protection includes such

mechanisms as secure configurations, data-at-rest encryption, and e-mail encryption. Although DON enterprise network resources are protected with zone 1 mechanisms, individual organizations and/or COI must implement their own zone 1 protection mechanisms.

E.4.4.1 Zone 1 Mechanism – Secure Operating Systems with Secure Configurations

The innermost layer of a defense in depth information protection approach is the computer operating system itself. Selection of operating systems for security features and the correct configuration of those features in an operational environment is critical. Guidance on the selection and configuration of securable operation systems is provided in Section 3.4.4.1 of the ITSG.

E.4.4.2 Zone 1 Mechanism – Content Security Checking

Viruses remain a significant problem for maintaining good information protection. DON information systems should use COTS virus protection software at servers and end user workstations. See Section 3.4.4.6 of the ITSG for further information.

E.4.4.3 Zone 1 Mechanism – E-Mail Encryption

COTS e-mail encryption can be used to provide user-to-user confidentiality, integrity, and authentication. Coupled with a properly managed Public Key Infrastructure (PKI), e-mail encryption can support CoI separation inside the enterprise, a region, a building, or even a workstation. See Section 3.4.4.3 of the ITSG for details.

E.4.4.4 Zone 1 Mechanism – World Wide Web (WWW) Encryption

COTS web encryption can be used to enforce user authentication and access control to information stored on a web server as well as to provide confidentiality, integrity, and authentication for information as it traverses a network. Coupled with a properly managed PKI, web encryption can support CoI separation inside the enterprise, a region, a building, or even a workstation. See Section 3.4.4.4 of the ITSG for details.

E.4.4.5 Zone 1 Mechanism – VPN Encryption

COTS VPN encryption can be used to provide confidentiality, integrity, and authentication for information as it traverses a network. Coupled with a properly managed PKI, VPN encryption can support CoI separation inside the enterprise, a region, or a building. See Section 3.4.4.5 of the ITSG for details. VPN protected modem access is acceptable.

E.4.4.6 Zone 1 Mechanism – Data at Rest Encryption

COTS data-at-rest encryption can be used to provide confidentiality and integrity for information stored on a workstation or server. See Section 3.4.4.2 of the ITSG for details.

E.5 Infrastructure Security Components

In addition to the information protection mechanisms located in the various zones, certain infrastructure components are required to build secure DON information systems. These components include a secure Domain Name System (DNS), Public Key Infrastructure (PKI), network management security support, and network security configuration management systems.

E.5.1 Infrastructure Mechanism – Secure Domain Name System

DNS is essential to the operation of most networked computers and applications. A secure DNS should be installed in ITSCs to reduce the potential of Naval systems being adversely affected by DNS attacks or misconfigurations of DNS servers outside the DON intranet.

E.5.2 Infrastructure Mechanism – Public Key Infrastructure

PKI is the primary security technology requirement. PKI technology will provide security at connection establishment time, but will not perform as an in-line ATM-level firewall. PKI, by itself, is not an adequate ATM-level security solution, but it is regarded as a prerequisite for a strong, flexible security solution.

E.5.3 Infrastructure Mechanism – Network Management Security

Remote management of network infrastructure components (e.g., routers, ATM switches, and SONET multiplexers) will be essential as DON networks are regionalized. If this remote management cannot be accomplished securely, these networks will be vulnerable to denial of service attacks. The simple network management protocol (SNMP) is often used to remotely monitor network components. It also can be used to remotely configure and control network components. Unfortunately, SNMP relies on a single unencrypted community string to control access to managed devices. As a result, SNMP is extremely vulnerable to community string interception and guessing attacks. The use of SNMP for remote monitoring is acceptable, but the use of SNMP for remote configuration and control should be avoided. At a minimum, user name and password should be used for access control to network components that must be remotely configured and controlled. The use of a COTS token-based access control system is highly encouraged.

E.5.4 Infrastructure Mechanisms – Network Security Configuration Management

As DON networks are regionalized and system administration and support functions are centralized to reduce TCO, it will become increasingly important to be able to quickly and automatically determine the security configuration of remote computers. This will allow system administrators at regional ITSCs to find security bugs in user workstations and ensure that appropriate steps are taken to eliminate these bugs. The use of COTS network vulnerability testing tools and security configuration checking tools by regional ITSC administrators is encouraged.

E.5.5 Infrastructure Mechanisms – ATM-level Security

If ATM service providers, including DISA, support only PVCs (they do not support SVCs) and Naval ATM switches connect directly to the service provider's switches, security would be ensured. In this case, the ATM WAN functions as a VPN. However, since the DISN ATM WAN and some commercial ATM service providers support SVCs, it is necessary to implement ATM-layer security to prevent unwanted access to the DON enterprise intranet. Therefore, access controls, at a minimum, are required. Also, because ATM is a connection-oriented protocol, once an ATM connection is established, no intermediate systems take part in filtering higher-layer

information. A virtual channel terminating on an end-system inside the DON intranet will bypass traditional IP firewalls. Thus, even with sufficient ATM access controls, there is no method to filter network layer (e.g., IP) data along the link once it has been established. Doing so would make them packet switches, not ATM switches. Therefore, it is not clear just how, or whether, it might be possible to implement firewalls in an ATM environment.

ATM-layer security is essential for the protection of the DON enterprise network. However, at present there is a clear lack of technical standards for ATM-layer security. The ATM Security Framework 1.0 is intended to provide a common language and structure for future technical standards, but it defines no implementable security protocols or algorithms. Commercial ATM-level security is still immature and where it exists, it tends to be vendor-specific. Because IP-layer attacks represent the greatest opportunity for would-be hackers, a lack of ATM-level security does not increase the overall security threat. This architecture identifies technologies that are expected to provide solutions to the security issues and will be a participant in DoD efforts to resolve ATM-layer security.

E.5.6 Infrastructure Mechanisms – IP Security

Network security above the ATM level includes IP security and application security. Although this architecture prescribes ATM to the campus and enables end-to-end (desktop-to-desktop) ATM, classical LANs and IP capabilities will continue to be supported. Therefore, enterprise-wide IP security components are required.

Volume I, Appendix F

ITI Technical Subteam Participants

ITI Architecture IPT Leader,

Don Endicott, SPAWAR Systems Center

ITI Architecture Technical Subteam

Technical Subteam Leader, Ron Broersma*, SPAWAR Systems Center

Wayne Beck, CNO N6

Rex Buddenberg*, Naval Postgraduate School, Monterey

Jim Coffman, NCTC

Basil Decina, NRL

Dale Edgeington*, BUMED

Gary Garris, NCTAMSLANT

Richard Glover, MARCORSYSCOM

Mike Harrison*, SPAWAR

Tony Ho, CINCPACFLT

Bob Johnson, CNO N46

Eric Markland, CINCLANTFLT

Dave Mihelcic*, NRL

LT Jim Mills, CINCPACFLT

Jack Murphy, DON CIO

Mike Randall*, NAVAIR

Tom Scruggs*, DON CIO

Merrill Witzel, SPAWAR/CINCLANTFLT

* Writing Team

This page intentionally left blank.

Department of the Navy Information Technology Infrastructure Architecture (ITIA)

Volume II *Enterprise Architecture* *Framework*

Department of the Navy Chief Information Officer (DON CIO)
Information Technology Infrastructure Integrated Product Team (ITI IPT)

Version 1.0 proposed
16 March 1999

This page intentionally left blank.

Introduction to Volume II

This publication is included as part of the Report of the Department of the Navy (DON) Chief Information Officer (CIO) Information Technology Infrastructure (ITI) Integrated Product Team (IPT). Volume II describes the Enterprise Architecture Framework (EAF) for the DON. This framework is being advanced under the leadership of the DON CIO and the CIO Board of Representatives and through the coordination of the Enterprise Architecture and Standards Competency Unit within the Office of the DON CIO.

This EAF is being introduced in conjunction with the results of the ITI IPT to provide the overall context for all enterprise architecture modeling activities. It positions the role of Information Systems and ITI within a broader strategic context related to the Revolution in Military Affairs and the Revolution in Business Affairs (RMA/RBA).

It is intended that this framework will assist the various planning teams across the DON by providing a unified and common set of reference models around which to plan and coordinate the complex changes involved in RMA/RBA.

It is intended that this framework will further evolve through effective application to various architecture initiatives in the Joint community as well as across all DON functional areas.

Volume II - The Enterprise Architecture Framework

The volume is organized into the following chapters:

Chapter 1 - The Enterprise Architecture Framework Overview

This chapter introduces the EAF. It discusses the purpose behind introducing a Department-wide architecture planning framework and presents some underlying principles to guide the development of this framework. The chapter introduces the overall framework and explains the purpose for each of the four primary views of the Enterprise Architecture.

Chapter 2 - The Mission View

This chapter describes the elements that comprise the Mission View. These are the elements that are used to construct strategic plans and derive the enterprise missions from analyzing the changing national security needs and other forces for change that are acting upon the enterprise. The critical role of Stakeholders, Partners, Suppliers, and Recruits is also included in this view.

Chapter 3 - The Operational View

This chapter describes the elements that comprise the Operational View. These are the elements used to construct operational capabilities in support of the Mission Essential Tasks and in relation to the required operational environments. This view includes Personnel, Platforms, Facilities, Equipment, Supplies, and Information.

Chapter 4 - The Systems View

Chapter 4 identifies the many different types of systems that are required to support the operational capabilities of the DON. These are presented in the EAF as sub-views. They include Weapons Systems, Information Systems, Other Special Purpose Systems, and their supporting Platform Architectures, Facility Architectures, Utility Architectures, and the Information Technology Infrastructure Architecture. The IT Infrastructure Architecture is the part of the Enterprise Architecture that was developed by the ITI IPT.

Chapter 5 - The Technical View

Chapter 5 presents the different categories of Technology Component Standards and Specifications relating to each of the systems' sub-views. For the IT Infrastructure, the corresponding standards are the Information Technology Component Standards as contained in the ITSG. The Standards Framework, as developed within the ITSG, is presented here to link this guidance document to the EAF.

Chapter 6 - Framework Implementation Considerations

The final chapter discusses the considerations for fully implementing the EAF. Topics include the need for developing reference models for the elements of each of the four views to create a standardized set of terms and definitions to be used across the enterprise. The need for creating a repository for enterprise architecture models and an associated tool set is also discussed. Chapter 6 also includes a section on how the EAF was applied to produce the work products of the ITI IPT.

Table of Contents

| | |
|---|------------|
| Introduction to Volume II..... | i |
| 1. The Enterprise Architecture Framework Overview..... | 1-1 |
| 1.1 Framework Principles | 1-1 |
| 1.2 Architecture Framework Terminology | 1-3 |
| 1.3 Framework Positioning..... | 1-4 |
| 1.4 C4ISR Framework Version 2..... | 1-5 |
| 1.5 Enterprise Architecture Framework – Four Views | 1-6 |
| 2. The Mission View | 2-1 |
| 2.1 Architecture Elements and Relationships | 2-1 |
| 2.2 Enterprise Missions | 2-2 |
| 2.3 Uses of the Mission View and Elements..... | 2-3 |
| 3. The Operational View | 3-1 |
| 3.1 Architectural Elements and Relationships | 3-1 |
| 3.2 Uses of the Operational View and Elements..... | 3-2 |
| 4. The Systems View..... | 4-1 |
| 4.1 Architectural Sub-Views..... | 4-1 |
| 4.2 The Information Systems Sub-View..... | 4-2 |
| 4.3 Uses of the Information Systems Sub-View..... | 4-3 |
| 4.4 The Information Technology Architecture Sub-View | 4-3 |
| 4.5 Uses of the IT Architecture Sub-View..... | 4-5 |
| 5. The Technical View | 5-1 |
| 5.1 Technical Component Sub Views..... | 5-1 |
| 5.2 Information Technology Standards Framework..... | 5-2 |
| 6. Framework Implementation Considerations..... | 6-1 |
| 6.1 EAF Implementation Critical Success Factors | 6-1 |
| 6.2 Using the Framework for the ITI IPT | 6-3 |
| 6.3 Conclusion | 6-8 |

This page intentionally left blank.

1. The Enterprise Architecture Framework Overview

The DON is involved in a major enterprise transformation as envisioned by the QDR notions of the RMA and the inter-related RBA. This revolution will affect most, if not all, of the operational areas of the Department. The complexity of planning and managing this transformation cannot be overstated.

The purpose for introducing this EAF is to assist all of the planning and implementation teams working on components of the RMA/RBA by providing a common underlying structure around which to model this transformation. This framework is intended to be comprehensive and address all of the inter-related components involved in planning and managing enterprise transformations.

The development of this EAF is a collaborative effort under the sponsorship of the DON CIO and the CIO Board of Representatives. The Enterprise Architecture and Standards Competency Unit of the Office of the DON CIO provides coordination of this initiative. This initial publication of the EAF is occurring in conjunction with the publication of the ITI IPT Report because this is the first of the enterprise-wide architecture initiatives to occur under the mandate of the DON CIO. A sub-team of the ITI IPT has been responsible for developing this initial version of the EAF.

As noted above, the EAF is intended to provide the structure and components for all areas of transformation planning and incorporate many disparate architecture initiatives from across Joint and Departmental planning areas. This EAF has built upon the existing C4ISR Version 2 Architecture Framework to ensure compatibility with this key Joint direction. It has extended this framework to be more broadly applicable to all functional areas and to tightly link investment planning in systems and infrastructure to the key strategic missions and transformation objectives of the Department.

The role of information technologies in supporting or enabling the intended transformation (Information Superiority, Network Centric Warfare, etc.) is strongly represented in this proposed EAF.

1.1 Framework Principles

In developing this Framework, the IPT defined a number of key principles that must be applied in creating an effective architecture planning capability to assist planners. These principles are described below along with their supporting rationale.

1.1.1 Enterprise Perspective

To achieve the stated strategic objectives of RMA/RBA, the DON must take an enterprise perspective in all planning of mission, operational, systems, and technology requirements and solutions.

Rationale:

- To fully explore opportunities for transformation, it is necessary to remove all existing functional, organizational, and geographic barriers.
- The impacts of information technology on how our military and business affairs will be conducted are far-reaching and require an enterprise view of how information can be better leveraged to support our primary mission objectives.
- The opportunities to leverage common solutions to mission and business requirements are significantly increased by combining the interests of all communities across the DON and by extending this analysis across Joint, Allied, other governmental agencies, and the private sector.

1.1.2 One Framework

To support enterprise oriented transformation planning, all of the various communities of interest within the DON must adopt a common framework with common references to architecture elements.

Rationale:

- A unified framework maximizes the opportunity to identify areas of commonality and opportunities for sharing in the development and use of common solutions.
- A unified framework supports the identification of interoperability requirements.
- A unified framework facilitates developing enterprise perspectives and “out-of-the-box” thinking.
- Using a common framework for tactical and non-tactical areas will assist in identifying opportunities to shift resources “from the tail to the tooth.”
- A unified framework facilitates the development of common planning and design methods and the ability to share and/or redistribute resources.

1.1.3 Fully Integrated

The framework must support modeling and planning of all required enterprise capabilities crossing traditional tactical and non-tactical functions and covering all operational environments.

Rationale:

- The framework must provide a uniform basis for assessing existing capabilities and prioritizing requirements to upgrade or improve areas of concern.
- Planning for the delivery of basic Infrastructure Services across the DON requires enterprise-wide coordination and synchronization.
- Although regionalization is concentrating control of IT resources, many communities of interest operate cross-regionally and need to integrate with the afloat and deployed commands, thereby requiring enterprise-wide views of current and planned capabilities on which to base their transformation strategies and plans.

1.1.4 Multiple Views

The framework must support the needs of all types of architects and planners and provide views of mission capabilities, operational capabilities, systems capabilities, and technology capabilities as well as the inter-relationships between these views.

Rationale:

- The planning of enterprise transformation occurs at a number of different levels, ranging from the strategic planning of enterprise missions down to the development and application of new technologies.
- Each of the levels involves different types of planners and architects with different experience and capabilities all working on inter-related aspects of the transformation plan.
- The framework must provide the top-down context for all of these initiatives such that strategic relevance and investment or resource allocation priorities can be linked.

1.1.5 Time-line Dependencies

To support current and opportunity assessment, alternative evaluation, target solution design, and migration planning, the framework must support the identification of time periods and interdependencies across implementation projects.

Rationale:

- Operational capabilities are dependent upon a number of factors aligning in time to ensure readiness to perform expected missions.
- These factors include the required number of suitably trained and organized personnel with appropriate platforms and/or facilities containing necessary equipment and supplies providing access to the required information systems.
- These systems require the appropriate IT infrastructure to be in place with adequate capacity, high availability, and supporting the necessary levels of security.
- The framework must support the ability to align these requirements and ensure synchronization across the various design and development, implementation, and training projects involved in creating the operational capability.

1.1.6 Supporting Toolset and Repository

A graphical architecture modeling tool and model repository is required to represent all EAF architecture models and elements with associated attributes and relationships.

Rationale:

- The framework is comprised of a number of architecture planning models and associated elements (objects). Most of these elements will have a large number of occurrences and many relationships. A tool and repository environment is essential to managing this complex information domain.
- Architecture drawings are graphical by nature. A graphically-oriented tool will allow appropriate “blueprints” to be developed and viewed and will standardize associated iconic representations.
- All architecture elements will have associated attributes that will be populated and updated as part of the model development or as the model is used to support assessment or planning.
- Architecture elements and models have many relationships that are developed or explored over the course of developing models, conducting assessments, and building plans. The tool must support the creation and viewing of these complex inter-relationships.

1.2 Architecture Framework Terminology

This section provides definitions of terms that are associated with the EAF. A key aspect of introducing an EAF is to provide a common means of referring to the components of the Enterprise. Here, we are practicing what we preach by normalizing our own use of terminology within the architecture community.

Architecture - the structure of “anything”; an image of the required or planned functionality.

Architecture Framework - the underlying or basic structure upon which other architectures are built.

Enterprise Architecture Framework - the underlying structure for planning the capabilities of an enterprise, encompassing all required functions and views.

Architecture View - a different perspective of the architecture provided for specific purposes and planners with different interests. The EAF identifies four views: Mission, Operational, Systems, and Technology.

Architecture Element - a primary component or enterprise “object” defined by the framework. Each element is a unique class (e.g. ship) with possible sub-classes (e.g., carrier, destroyer, etc.) with associated attributes (e.g., tonnage, maximum speed, date commissioned, etc.) with definable relationships (e.g., member of battlegroup X, SIPRNET-enabled, embarked commands, etc.)

Architecture Model - Any representation using combinations of architectural elements and possibly including their attributes and relationships. These can be diagrams showing groups of elements, connectivity or flow between elements, tables showing relationships, or reports containing diagrams, tables, and textual descriptions (e.g., an information flow model showing how information is created and used by a related operational functions). These models can represent different periods of time from current or baseline assessments through planned implementations.

Reference Model - Identification of the set of occurrences for Architectural Elements and their primary relationships in advance of their use. This provides a common means of identifying and defining elements for use across the enterprise (e.g., the Universal Joint Task List -UJTL for joint operational functions and tasks).

Infrastructure Service - Recognition of a common service or capability that is widely required across the enterprise or by large numbers of users such that it can be planned and operated as a “common enterprise utility” (e.g., communication networks).

Template - A pre-planned architecture model providing solution guidance for a recognized common or repeatable set of requirements (e.g., metropolitan area network template, service center management and organization template).

1.3 Framework Positioning

Based on these definitions, the Architecture Framework can be positioned within the process for planning for transformational change. Figure 1-1 summarizes this positioning for an architecture framework.

An Architecture Framework provides the structure required to translate a vision for a renewed enterprise capability into a well-planned and managed set of implementation projects.

A framework is really a model of models (or a meta-model) because it describes and positions a number of different architectural views that are made up of different architecture elements. These views and elements are then used to construct different architecture models that assist planners and designers in making key decisions regarding opportunities for transforming the enterprise.

In many areas, these models will identify common requirements from which templates can be developed to leverage reusable solutions.

By providing visualization (a blueprint) of the desired change and its inter-relationships with other elements, the models and templates help the planners focus investments in the areas of greatest benefit to the enterprise.

By effectively applying this managed architecture approach, it is possible to prioritize and sequence implementations in accordance with overall enterprise strategic goals.

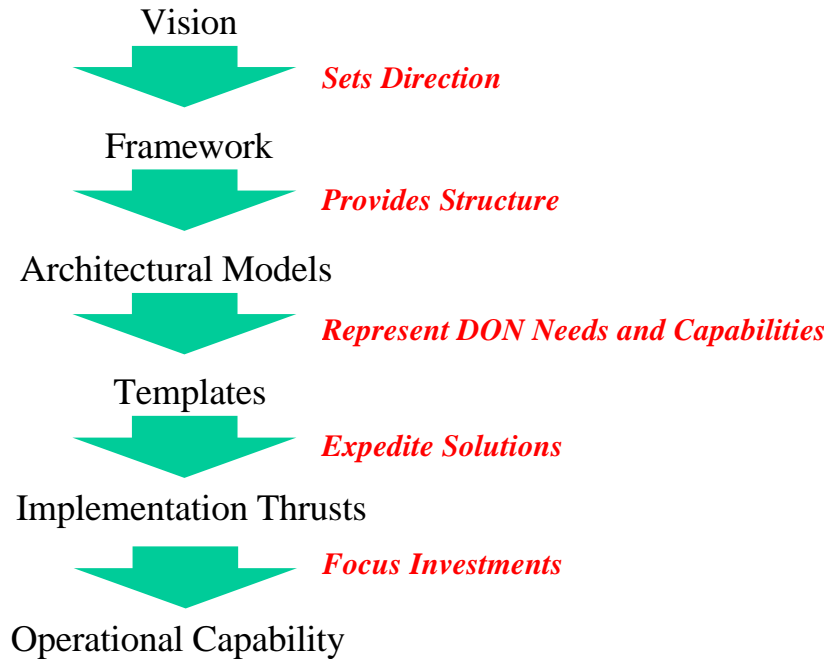


Figure 1-1. Positioning the Role of the Framework

1.4 C4ISR Framework Version 2

Figure 1-2 depicts the current version of the C4ISR Architecture Framework. There are three different Architecture Views - Operational, Systems, and Technical - and a corresponding set of inter-relationships.

This framework is the basis for much architecture planning work underway across the Joint Services in the C4ISR Community. This EAF was built upon this base.

This existing framework is now heavily populated with reference models for many of the elements and relationships. It is rightly focused on the needs of the warfighter and supports the integrated requirements for operational, systems, and technology planning

The EAF has built on this foundation and extended this framework in the following areas:

1. Added the “Mission View” to provide a structured approach to representing strategic operational requirements as driven from RMA/RBA and the corresponding threat analyses
2. Extended the functionality to include all required enterprise functions (operational, support, and planning) and all operational environments
3. Expanded the “Systems View” to include all types of “Systems” (Weapons, Platforms, Information, etc.) and added the notion of IT Infrastructure to the Systems View to address the need for templates based on the established Technical Standards (JTA and ITSG)
4. Positioned each of the four views in the flow of requirements and capabilities to show the inter-dependencies and emphasize the need for alignment

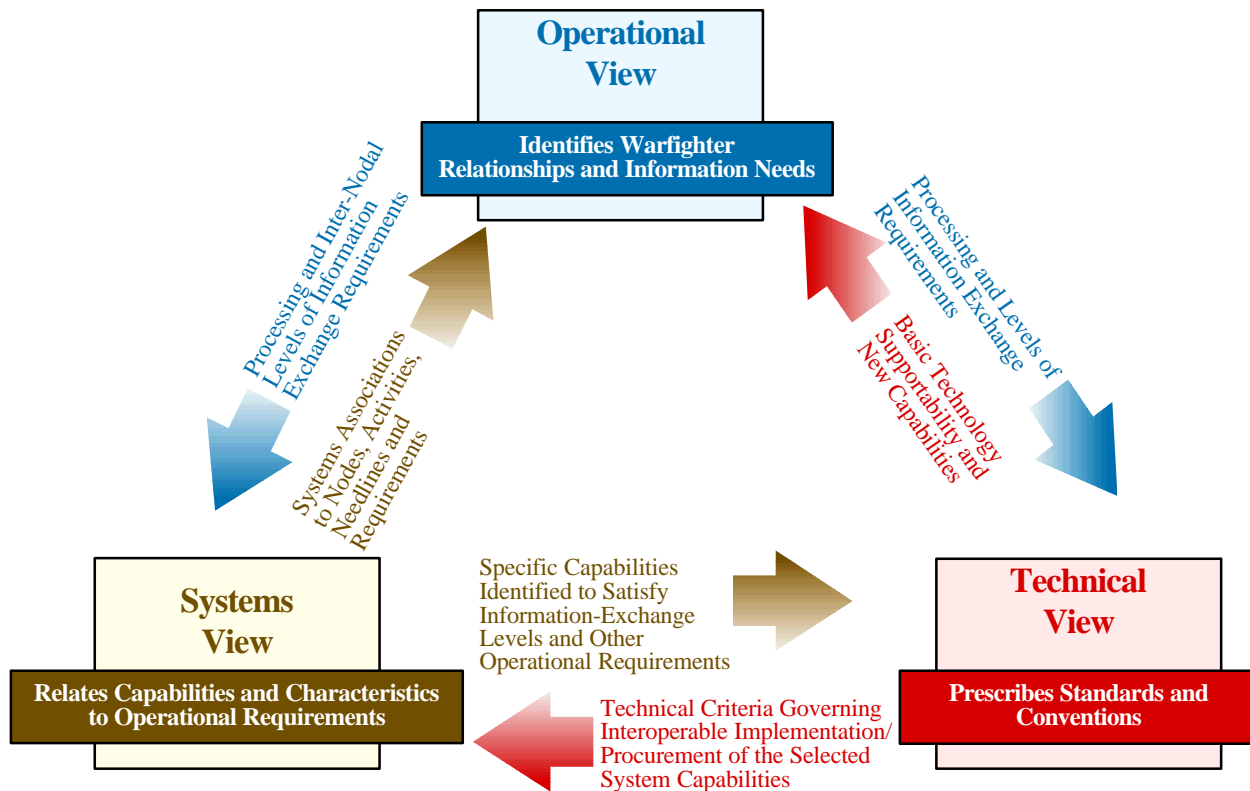


Figure 1-2. The C4ISR V2 Architecture Framework

1.5 Enterprise Architecture Framework – Four Views

Figure 1-3 shows the four top-level views of the proposed DON EAF. The Mission View has been added to provide additional structure pertaining to strategic mission requirements.

The Operational View has been expanded to reflect all functional requirements of the Enterprise (tactical and non-tactical) in a common structure with associated environments.

The Systems View now contains all types of systems and the Technical View provides the corresponding technology standards for the various systems components.

All of these four views have key inter-relationships. Requirements flow from top to bottom and capabilities flow up to address those requirements.

The goals of the framework include:

1. Achieving complete alignment of capabilities with requirements across these different views. The Mission View aligns the determination of mission areas and associated funding with the needs for Naval military services and develops associated operational requirements. The Operational View responds to these requirements by delivering the appropriate operational capabilities while at the same time imposing requirements on the Systems View and the underlying technologies. These views respond with the corresponding capabilities.
2. Establishing the strategic priorities for funding and then tracing these priorities through the views to determine the inter-dependencies on capabilities (e.g., how does improving a key mission area relate

to new operational capabilities that depend on what systems and infrastructure that are built upon what technologies).

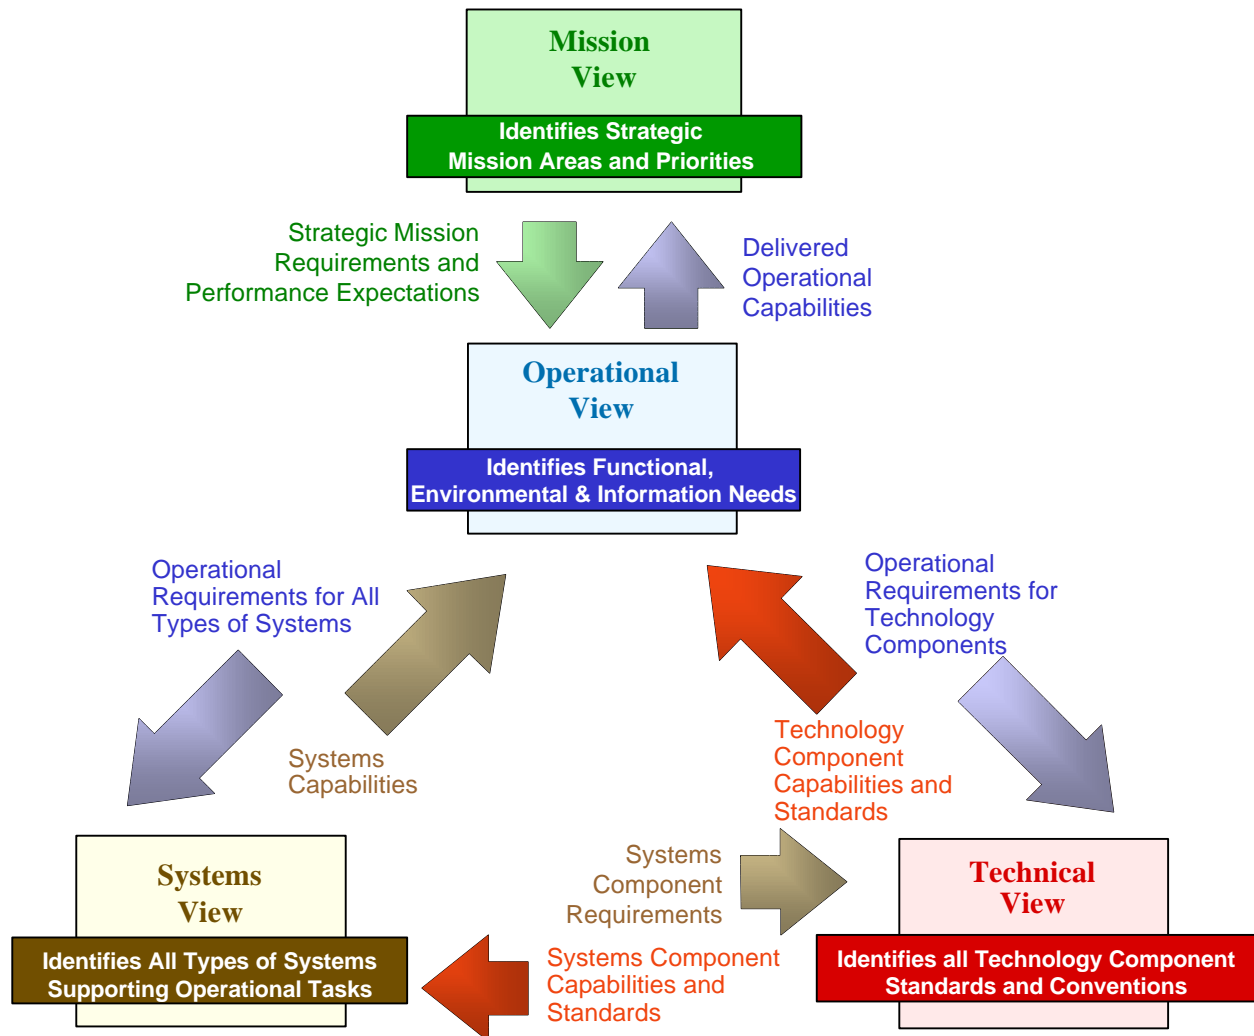
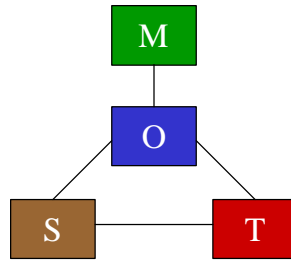


Figure 1-3. The Enterprise Architecture Framework Showing Four Views

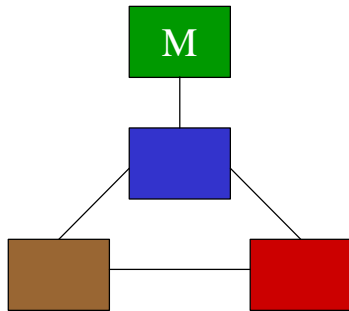
3. Developing effective governance and performance management capabilities. What are the fundamental changes in organization/command structures and accountabilities? What performance measures are needed to motivate appropriate behavior?

To achieve these goals, it is necessary to develop a comprehensive framework that includes all of the required planning elements and their inter-relationships. This top-level representation of the EAF is referred to as the "MOST" model using the following icon. The subsequent four chapters describe the underlying components and relationships for each of these four views.

Draft Working Papers of the ITI IPT



2. The Mission View



The Mission View represents the highest level of the EAF. It is used to model the elements and relationships involved in developing strategic plans for the DON.

This chapter presents the architecture elements and relationships that are involved in the Mission View. It is intended as an introduction to the use of architecture models in supporting strategic and mission planning. It also provides the critical link to the interests of the key stakeholders of the DON as represented by the U.S. Government and the many influences that these stakeholders have on the strategies and missions of the DON.

2.1 Architecture Elements and Relationships

The Mission View is comprised of a number of architecture elements that are used in developing strategic plans and setting the goals and priorities of the DON. Figure 2-1 presents these elements and the key relationships between them.

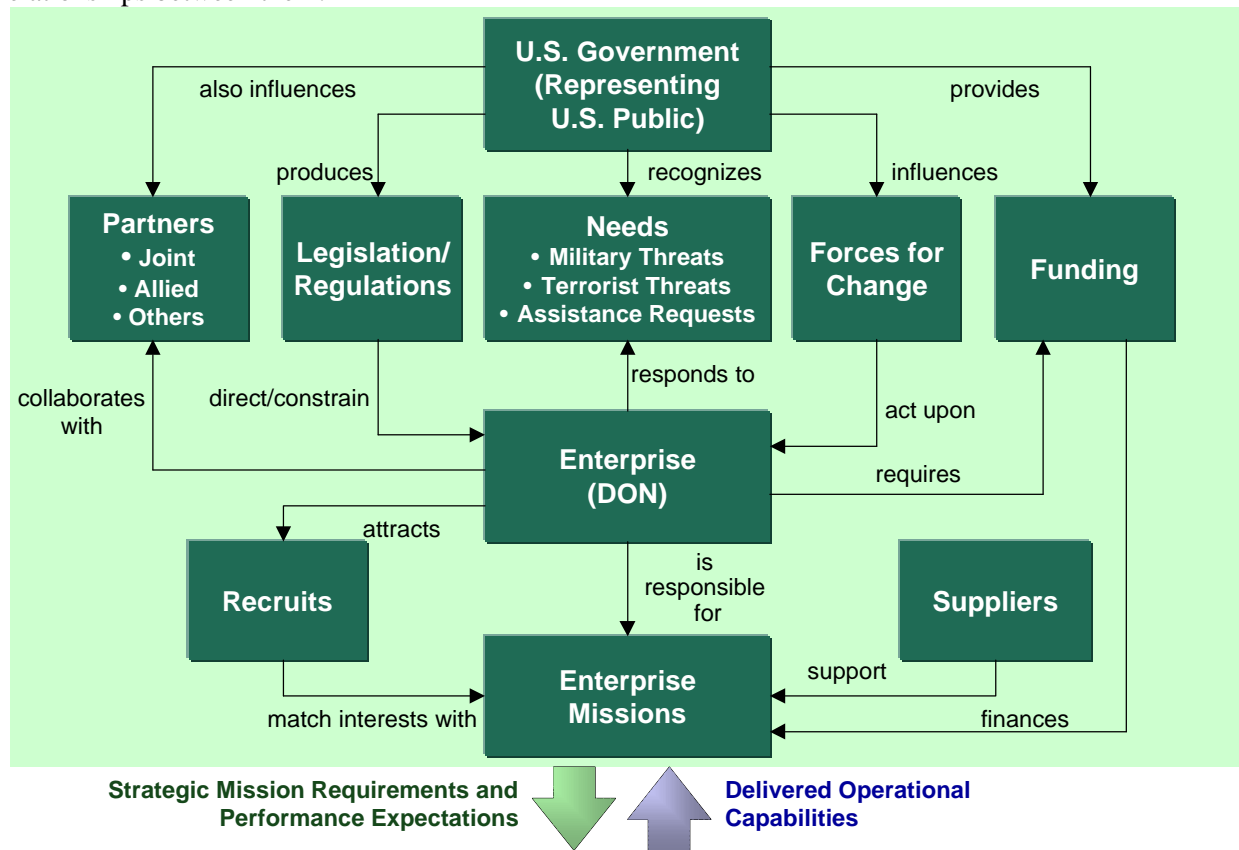


Figure 2-1. The Architecture Elements and Relationships in the Mission View

There are ten (10) architectural elements that comprise the Mission View of the DON EAF. The following paragraphs provide a brief overview of these elements.

The stakeholders for the DON are the people and interests of the USA as represented by **the Government of the U.S.** with the elected bodies of the President, Congress, and the Department of Defense. In their governing role, these bodies and agency also oversee the Army, Air Force, Coast Guard, and other **partners** with whom the DON collaborates in developing integrated national defense strategies and in conducting missions.

Central to the development of defense strategies is the identification and evaluation of **Needs**, including combating threats with force or deterrents and meeting demands for other military assistance (peacekeeping, disaster assistance, etc.). The determination of the role for DON in addressing these needs is central to the development of strategy.

The core element produced by this analysis is the set of **Enterprise Missions** that define the role of the DON in relation to other services and agencies involved in National Defense Programs. These missions form the basis for determining **Funding** requirements and priority allocations. The alignment of funding with the Naval strategic requirements and the resulting operational capabilities is a key relationship of the DON with its stakeholders.

The Government also produces **Legislation & Regulations** which impose constraints and directives on the **Enterprise** (DON). They also influence the **Forces for Change** that are acting upon the DON (e.g., budget constraints, acquisition reform) along with many other sources (e.g., new technologies) which much be addressed in strategy development.

Strategies related to **Recruiting** requirements and effective use of **Suppliers** are also involved in finalizing mission strategies.

2.2 Enterprise Missions

The framework is used to identify architecture elements and relationships. These elements take on full meaning once they have been “populated” with the complete set of current or planned “instances” to form a reference model for all planners involved with this element. Figure 2-2 provides an example of the reference model for Enterprise Missions.

These missions were extracted from published strategy and vision documents and logically consolidated. Under the overall goal of providing full spectrum dominance, the model shows nine (9) primary missions. Within these primary missions, there are fifty-one (51) secondary missions identified.

This provides the structure to clearly classify and define these missions. The inter-dependencies amongst these missions can be traced back to the needs of combating various types of threats, providing deterrents, and offering military assistance. This analysis, along with other strategic considerations, can be used to help prioritize these mission areas for both funding purposes and to focus readiness and renewal activity.

Each of these strategic missions can be linked to the Operational View. Each imposes operational requirements, some of which are specific to the mission, but many of which are common across some sub-set or all missions.

As we proceed into the Operational View, it is important to think about mission requirements in terms of pulling through Mission Essential Tasks which can be shared across mission areas. In this way, we avoid re-creating organizational “stove-pipes” around missions.



Figure 2-2. Reference Model Showing Primary and Secondary Enterprise Missions

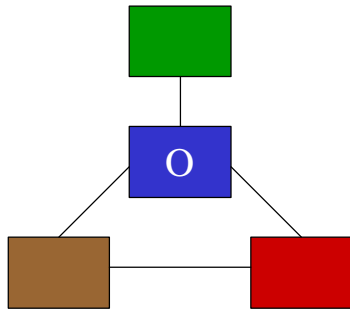
2.3 Uses of the Mission View and Elements

As illustrated above with the “Enterprise Mission” element, the framework provides the structure for populating architecture elements and tracing relationships or inter-dependencies. The elements within the Mission View can be used to construct many different types of models and perform evaluations and analyses. Some of these are listed below.

1. National Security Threat Analyses vs. Mission Capabilities
2. Strategic Impact Assessments of Forces for Change
3. Weighting of Enterprise Goals and Strategic Drivers to Mission Areas
4. Inter-dependencies of Enterprise Missions (or Sub-Missions) to Needs
5. Distribution of Funding by Mission/Sub-Mission
6. Mission Partnering Opportunity Analyses
7. Pending Legislation Impact Analysis
8. Supplier Dependencies - Risk/Opportunity Analysis
9. High-Level Funding Impact Assessments
10. High-Level Enterprise Readiness and Performance Assessments

As part of the implementation plan for the EAF, it is intended to connect with the planning communities involved in these models and analyses to develop meaningful examples of how the Mission View can be fully used to support strategic planning and provide the all-important context for operational planning. See Chapter 6 for information regarding these implementation plans.

3. The Operational View



The Operational View is positioned in the heart of the EAF. As the name implies, this view is very physical in nature and is used to describe the operational requirements and capabilities as dictated by the Strategic Mission requirements defined in the Mission View.

This Operational View sets the context for defining the requirements for all types of systems and technologies as they relate to the needs of the task, the people performing those tasks, and the environment in which they are performed.

This view is used to identify the operational capabilities which define “readiness” to fulfill the strategic mission requirements.

3.1 Architectural Elements and Relationships

The Operational View is comprised of a number of architecture elements that are used in developing mission tasking, describing operational capabilities, and identifying operational environments. Figure 3-1 presents these elements and the key relationships between them.

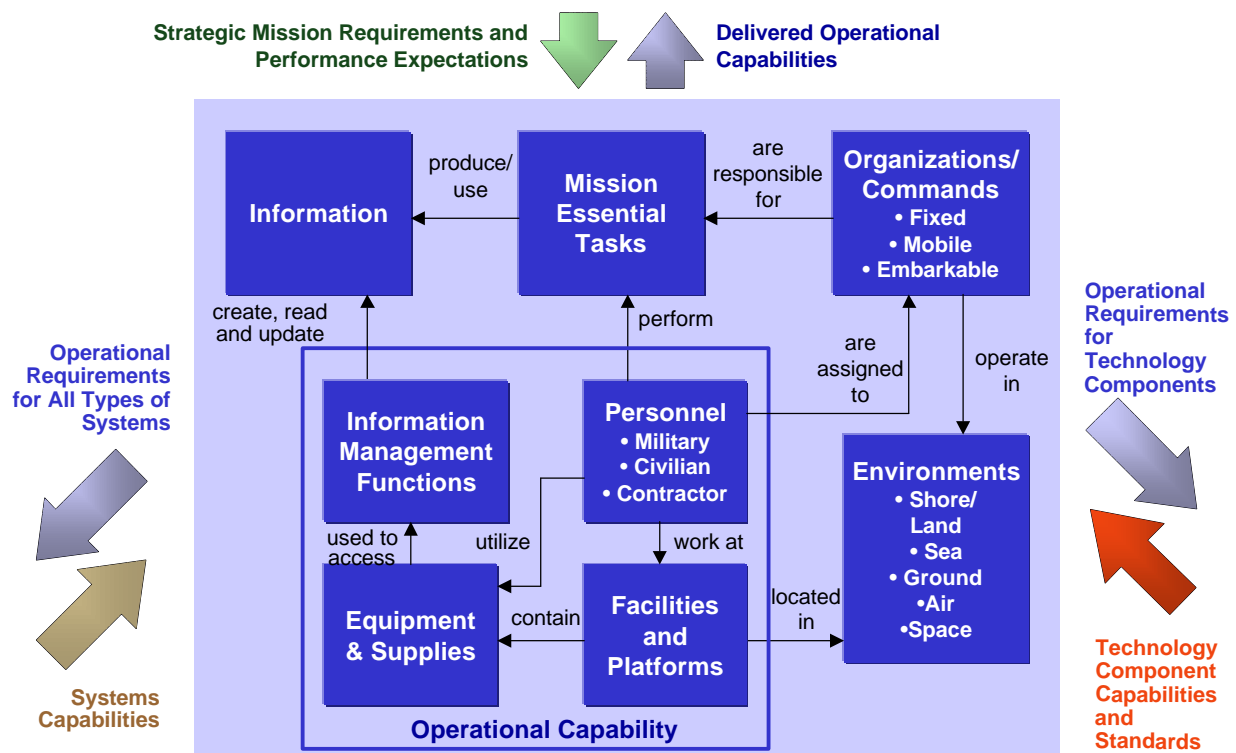


Figure 3-1. The Operational View Showing Architecture Elements and Relationships

The **Operational View** is comprised of a number of architecture elements that collectively are used to describe the overall operational capabilities that are required to respond to the Strategic Mission Requirements.

Mission Essential Tasks (METs) provide the means for modeling the functional requirements of the Enterprise to respond to the stated missions. For the DON, they are the Universal Navy Task List (UNTL), which is based on the UJTL. These tasks are grouped into four layers with multiple levels of decomposition. These tasks can be mapped to the **Organization and Command** Structure that is similarly structured as a hierarchical decomposition structure. Note that the UNTL was intentionally very mission-focused and must be expanded to accommodate many additional enterprise planning and support functions.

Depending on the METs to be performed, a command can be classified as fixed, mobile, or embarkable depending on its relationship to operating **Environments**. A **fixed** unit is assigned to a permanent workplace (in a facility or on a platform). A **mobile** unit moves within its associated environment (e.g., an IPT), and an **embarkable** unit is one that moves within and across operating environments (e.g., Marine Corps and Special Operations units).

The **Information** element identifies and structures the information resources of the enterprise. It provides the basis for defining information flows between METs.

The **Operational Capabilities** required to meet the needs of the METs are defined as combinations of **Personnel** with the appropriate skills supported by their work environments including necessary **Equipment and Supplies** and the related **Facilities and/or Platforms**. Some of that equipment (e.g., PCs, radios, and phones) is used to access the **Information Management Functions** that are used to manage the information requirements of the MET. These operational capabilities set the context for engineering systems solutions as represented by the Systems View. They also impose certain requirements on the underlying technologies that comprise these systems as represented in the Technical View.

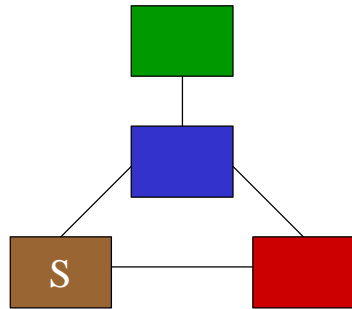
3.2 Uses of the Operational View and Elements

The following models and analyses based on the Operational View can represent **current, proposed, targeted, and migration states where applicable**:

1. Mission Essential Task Lists
2. Work Flow Models
3. Information Flow Models
4. Organizational Accountability Models
5. Resource/Cost Allocation Models
6. Operational Performance Assessment Analyses
7. Operational Improvement Business Case Analyses
8. Systems Requirements Analyses
9. IT Infrastructure Requirements Analyses
10. Funding and Project Impact Analyses

As part of the implementation plan for the EAF, it is intended to connect with the planning communities involved in these models and analyses. This will lead to development of examples of how the Operational View can be fully used to describe the required operational capabilities and provide the context for systems and technology planning. See Chapter 6 for information regarding these implementation plans.

4. The Systems View



The Systems View of the EAF is used to represent the many different types of systems and their components that are required to support the operational capabilities of the DON. This is the “engineering view” in which complex systems are put together from the standardized components as decreed in the Technical View and must all work together to meet the complex needs of the user in their operational environment.

The Systems View has combined all of these engineered systems under one view, including the weapons and platforms of the warfighter, the sensor systems that provide the battlespace information, the shore-based facilities and all of the utility and information technology infrastructure that is required to support these other systems.

4.1 Architectural Sub-Views

The Systems View is comprised of a number of different architectural sub-views which each represent a major class of systems. These are shown in Figure 4-1.

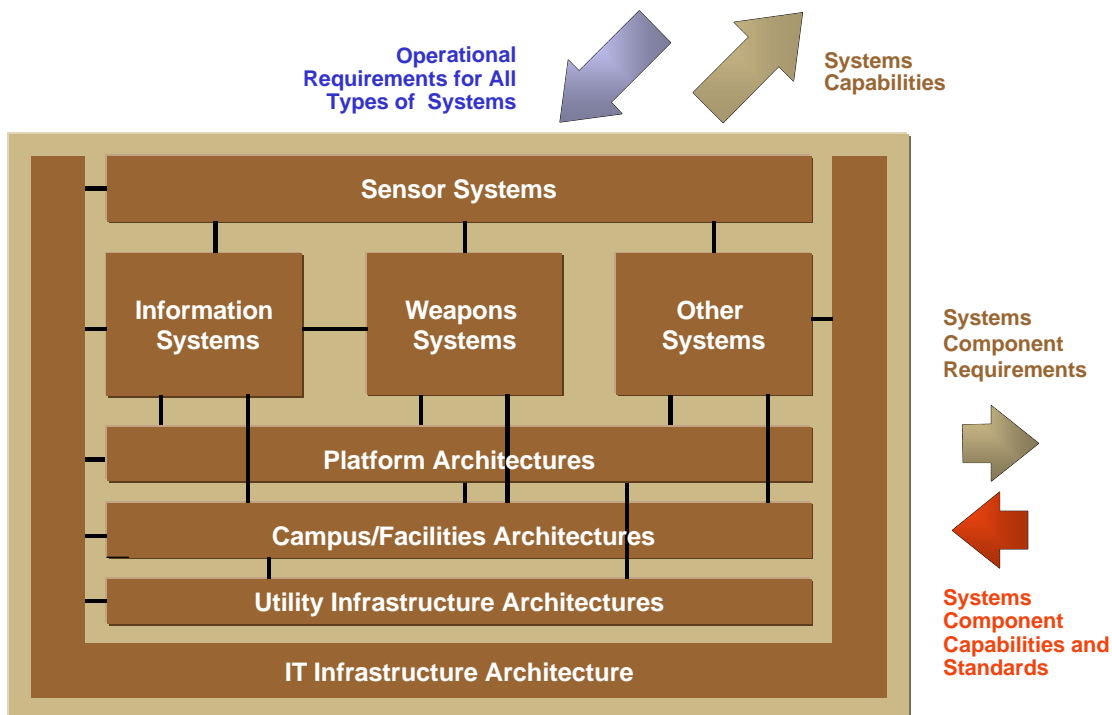


Figure 4-1. System Architecture Sub-Views

Note that these sub-views are still very high-level and may share many common components. They are used here to categorize major types of systems initiatives and how they come together to form a complete operational system for the warfighter or other types of users.

Sensor Systems provide the overall surveillance and data collection capability required to support the planners, commanders, and warfighters. Collectively, various sensor systems combine to produce information “grids” required by various user communities.

Information Systems include all types of information processing and management applications. These are generally specific to Communities of Interest or are common across the Enterprise.

Weapon Systems are a distinct class of systems including the propulsion, guidance, and payload capabilities, all of which may imbed information technologies as part of their respective control systems.

Other Systems include simulators, trainers, robots, materials handling systems, and other sensor-based and/or real-time systems with special user interfaces.

Platforms, including ships, planes, spacecraft, amphibious units, and vehicles, are also recognized as a distinct class of systems with their respective architectures. **Campuses and Facilities** are also viewed as “systems” with their own classes, components, and relationships. The **Utility Systems**, including power, HVAC, water, and sewage, are a sub-view which supports platforms and campuses/facilities.

The **Information Technology Architecture** supports all of the above sub-systems as well as directly interfaces with the various user communities in the operating areas. It is further described in Section 4.5 of this chapter. The ITI IPT used this sub-view to establish a set of enterprise planning templates for the integrated network infrastructure for the DON.

4.2 The Information Systems Sub-View

With the formation of the DON CIO Enterprise Systems IPT occurring at the time of the completion of the ITI IPT, it is important to connect this initiative with the EAF. For this reason, we are including the Information Systems Sub-View at this time. Figure 4-2 shows the layers that comprise this sub-view.

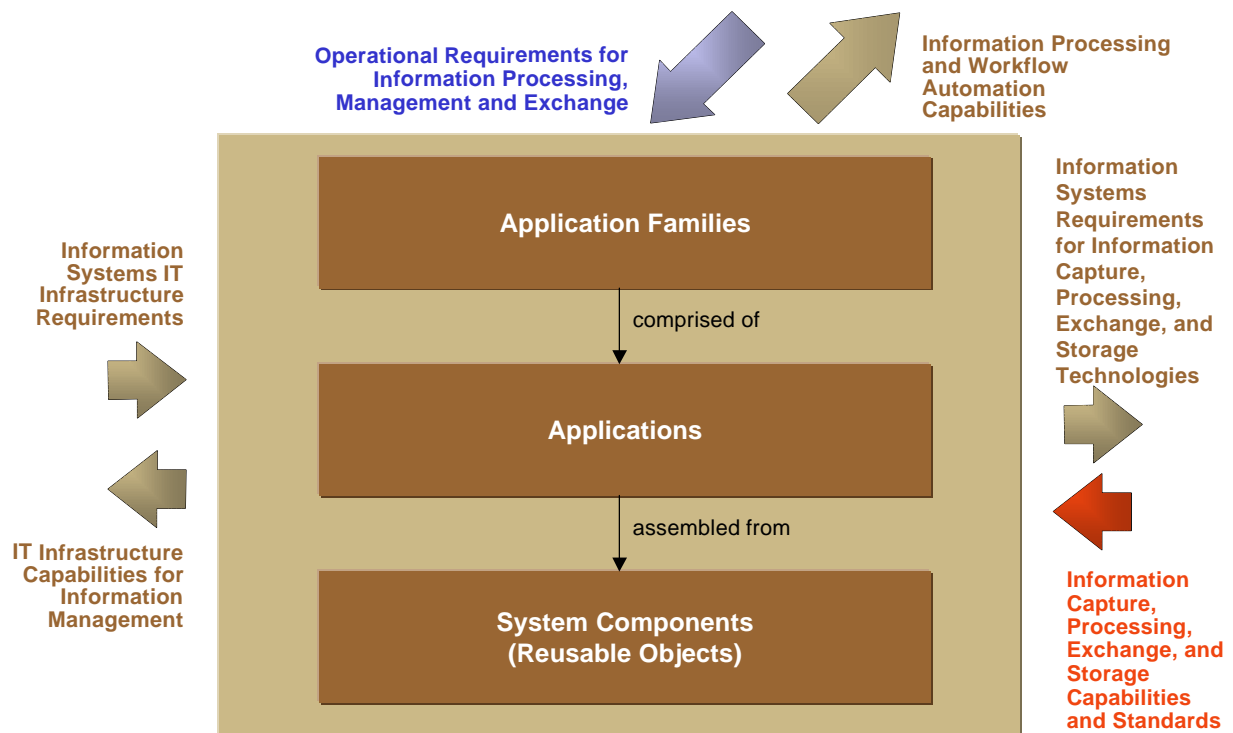


Figure 4-2. Information Systems Sub-View Showing the Three Layers

The **Information Systems Sub-View** includes all of the enterprise IT applications. These applications are related to the information capture, processing, and management requirements derived from the operational capabilities represented by METs and their information needs.

Information systems are first clustered into **Application Families** (e.g., GCCS, financial systems, and H.R. systems). These application families are then further broken down to identify specific **Applications**. Applications will typically have relationships with other applications within a family and may also share or exchange information with applications outside its family.

A key architectural construct in systems development today involves identifying **Application Components** or **“Objects.”** These are the building blocks for assembling application capability. In the object world, these capabilities are referred to as **“Use Cases,”** which is a powerful way of linking to the Operations View. A use case can be considered an application requirement invoked when a user requests a function from an application or an event occurs that triggers an application.

Application components are shared across applications and lead to the creation and support of reusable component libraries to aid in the productivity and standardization of application systems development.

This area of the framework is intended to support the DON CIO Enterprise Systems IPT in establishing enterprise guidelines for shared systems. As well, the Information Systems Architecture will assist systems planners and architects in designing future applications solutions that are consistent with operational capabilities and priorities.

4.3 Uses of the Information Systems Sub-View

As with the other views, there are a number of different types of models and analyses that are supported by the Information Systems Sub-View. These include:

1. Relationship of Application Families and Applications to Mission Essential Tasks
2. Relationship of Application Families and Applications to Information
3. Relationship of Applications/Tools to Organizations
4. Inter-system Information Flow Analysis
5. Systems Architecture Models
6. Reusable System Component Libraries and Usage Analyses
7. Y2K Compliance Analyses
8. Preparation Analyses for introduction of Object Management

As part of the implementation plan for the EAF, it is intended to connect with the Enterprise Systems IPT and the systems planning communities involved in these types of models and analyses to develop meaningful examples of how the Information Systems Sub-View can be fully used. See Chapter 6 for information regarding these implementation plans.

4.4 The IT Infrastructure Architecture Sub-View

The primary underlying “system” for the DON in planning for RMA/RBA is the enterprise IT Infrastructure that supports all of the required information systems. The IT Infrastructure Architecture Sub-View provides the framework for developing this infrastructure and recognizing the dependency on implementing this capability to achieve the necessary inter-operability across METs.

Figure 4-3 shows the architecture components within the IT Architecture.

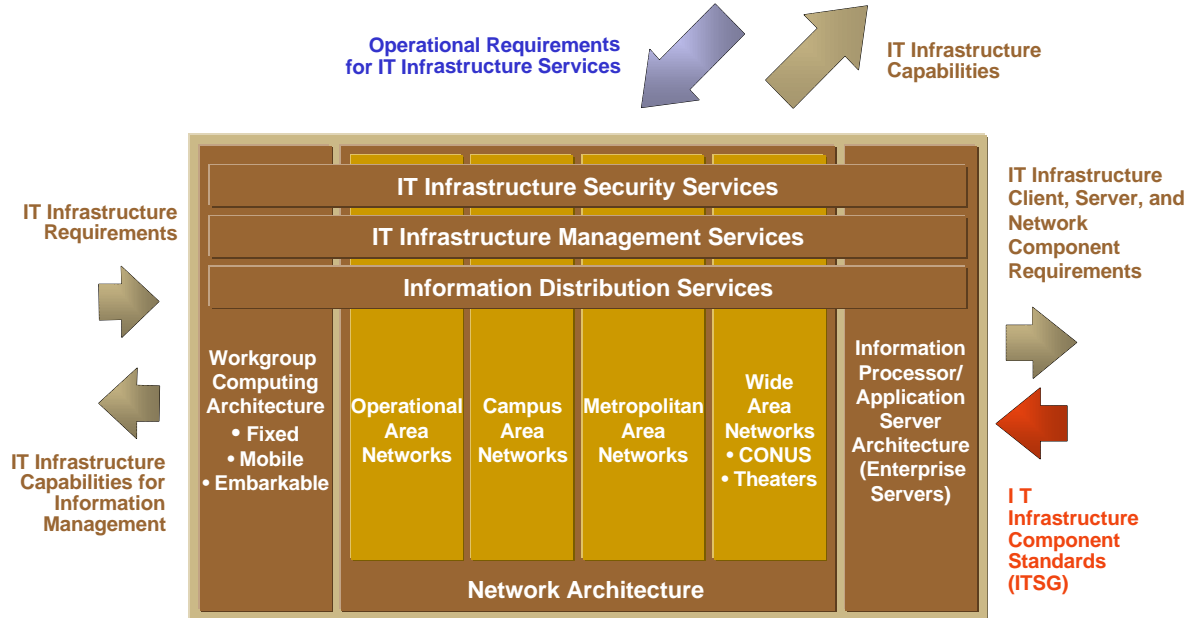


Figure 4-3. The IT Architecture Sub-View Showing Component Architectures

The IT Architecture contains all of the common information access, management, and exchanges services required by information systems and users of information technologies.

The **three underlying component architectures** are:

The **Workgroup Computing Architecture** that addresses the user access devices (PCs, telephones, radios, etc) various display devices, LANs, and workgroup servers. These architectures must support fixed, mobile, and embarkable implementations;

The **Network Architecture** that addresses the connectivity between workgroup devices and servers across the four levels of networks (operational, campus, metropolitan, and wide area) supporting ashore, afloat, and expeditionary communities; and

The **Server Architecture** that addresses the various information and applications processing requirements supported by the ITI.

This network computing capability is then over-laid with three **end-to-end infrastructure services**:

Information Distribution Services that provide for various types of information exchange and communication services across the ITI;

IT Infrastructure Management Services that provide performance and service level management capabilities plus other operational services; and

IT Infrastructure Security Services that provide for the stringent requirements for controlled access, information protection, and infrastructure management protection across the ITI for the various levels of security.

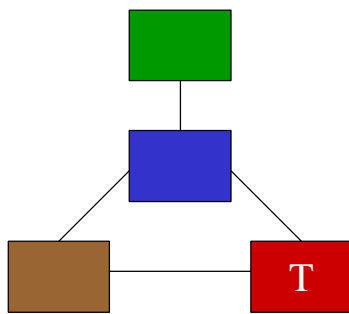
4.5 Uses of the IT Architecture Sub-View

The following list describes some of the uses for this sub-view. Note that the ITI IPT and other current planning initiatives are using these.

1. Development of Architecture Templates for Campus and Metropolitan Networks - Current ITI IPT mandate.
2. Development of Wide Area Network Connectivity Strategies and Plans - Current ITI IPT Mandate.
3. Development of Network Security and Network Service Architecture overlays - Current ITI IPT mandate.
4. Development of Network Management Center Architecture - ITSC Architectures being refined by ITI IPT.
5. Linkage with ITI Mission Essential Tasks to address Governance and Management Requirements - ITI IPT deliverable.
6. Development of Workgroup and Server Architecture Templates – future undertaking.
7. Development and Maintenance of Current and Planned ITI capabilities (BAM etc.)

The components of these various IT architecture templates are further represented in the Technical View. This provides the linkage of the ITI Architecture to the Information Technology Standards Guide (ITSG) as described in the next chapter.

5. The Technical View



The Technical View of the EAF provides the structure for organizing the complexity of standards and specifications associated with the components of the various types of systems represented in the Systems View. This is the “design” view of the framework in which individual building block specifications and their interfaces are determined.

The intention here is to rationalize the need for components and reduce unnecessary diversity. Although the special operational requirements for military applications of technology must be met, it is important to recognize the economic value and timing considerations involved in applying Commercial off the Shelf (COTS) and Government off the Shelf

(GOTS) technologies. This framework is also helpful in focusing research and development (R&D) efforts and evaluating emerging technologies.

5.1 Technical Component Sub-Views

The Technical View is comprised of a number of technology component areas around which standards and specifications are developed and maintained. Figure 5-1 shows these components:

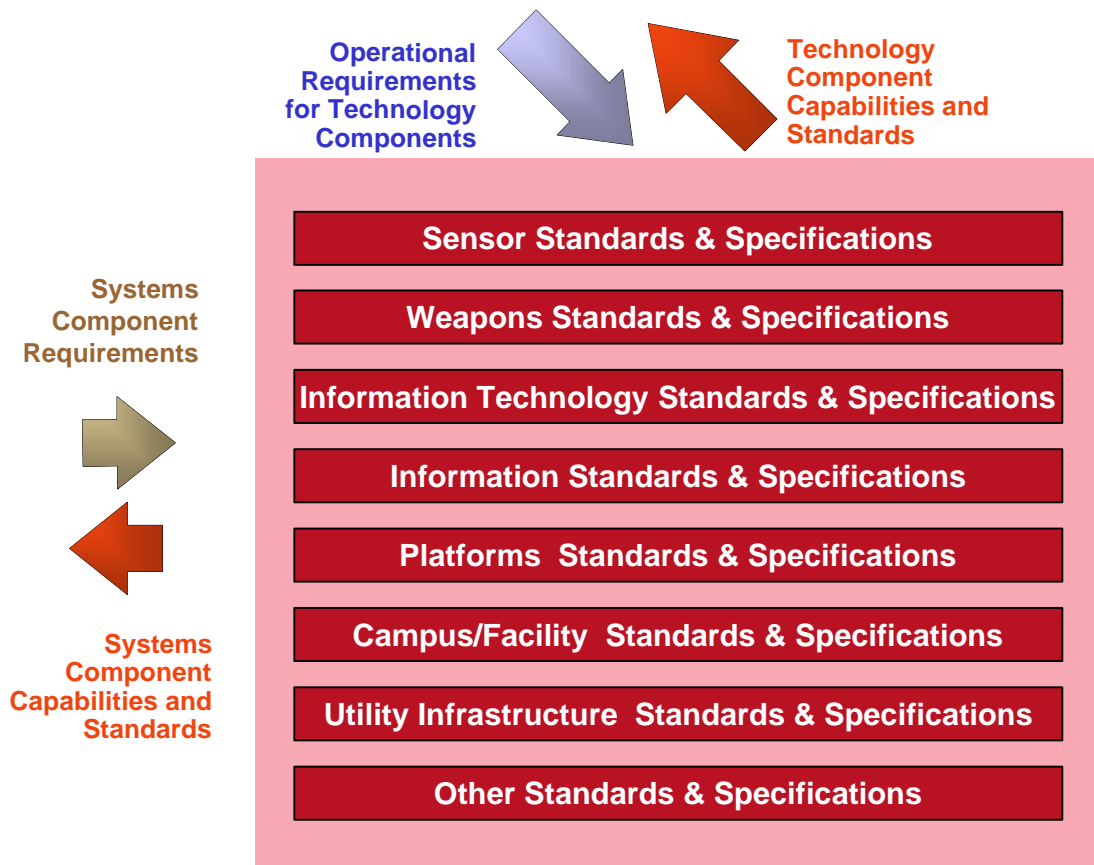


Figure 5-1. The Technical View Showing Components

These various components are extracted from the Systems View. The identification of components and their associated standards needs to be pulled together from the many sources to produce an enterprise-wide perspective on these technologies.

5.2 Information Technology Standards Framework

The initial version of the Enterprise Information Technology Standards was published in early 1998. The framework that was used to categorize these information technology components is shown in Figure 5-2.

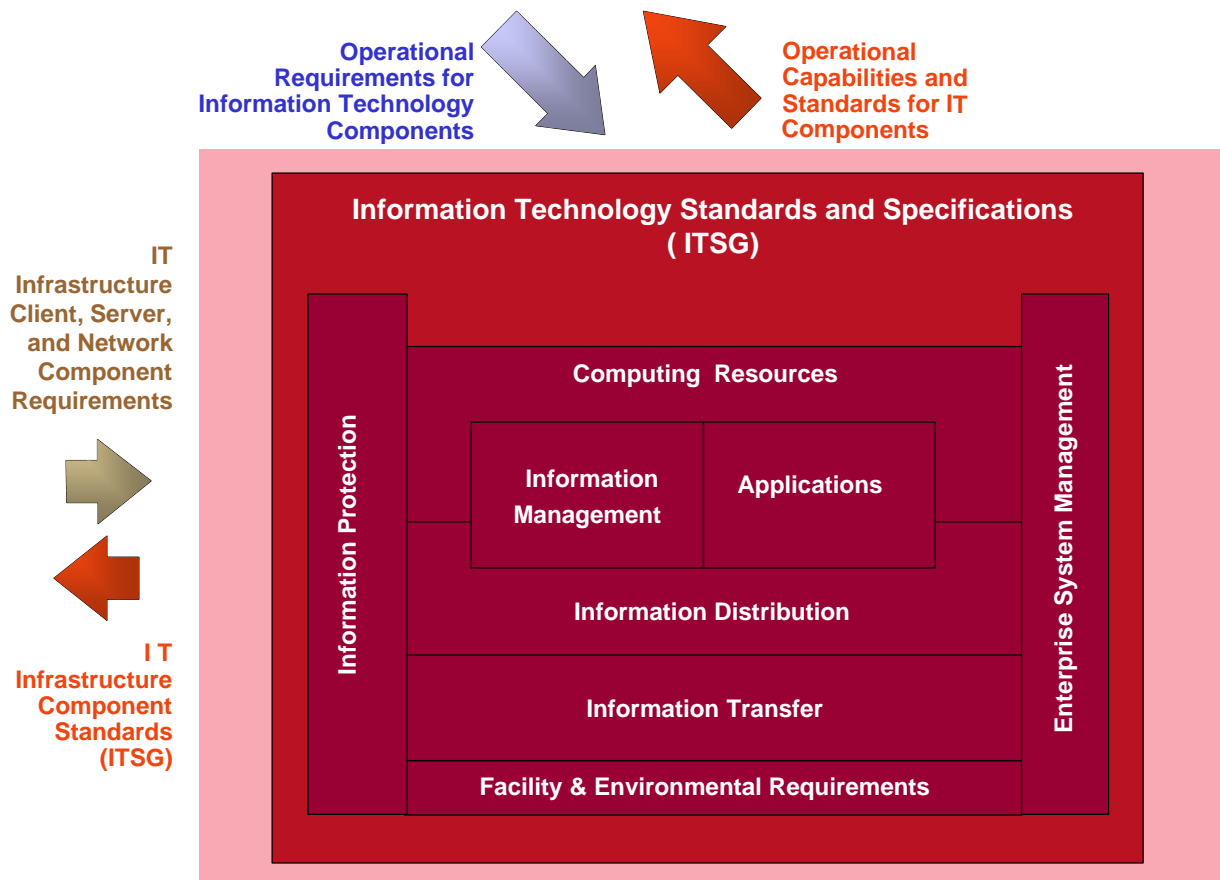


Figure 5-2. The Information Technology Standards Framework

These standards are linked to the ITI templates to add further definition to the ITI architecture for design and acquisition purposes.

These standards and guidelines must be maintained to remain current with information technology advances and market adoption of open systems and de facto standards.

6. Framework Implementation Considerations

The preceding five chapters have introduced the initial version of the DON EAF and described the components for each of the four architectural views - Mission, Operational, Systems, and Technical.

This chapter presents a number of key considerations regarding the implementation of the DON EAF and describes some short-term actions that are underway to capture existing enterprise reference models. In addition, it is important to provide a working prototype of an enterprise architecture modeling tool and repository. It is also important to link with key projects to support framework adoption and measure its effectiveness.

6.1 EAF Implementation Critical Success Factors

The publishing of this initial version of the DON EAF represents an important first step in establishing the means to coordinate and integrate the many planning activities that are underway across the DON. There are a number of follow-on activities that will continue to require collaboration and support from many different commands across the DON. The following Critical Success Factors have been identified for the implementation of the EAF.

6.1.1 Effective Coordination of EAF Evolution

The EAF requires on-going leadership and coordination to oversee its further development and to support its implementation across the DON. As experience is gained, updates will be required to the structures, the reference models, and the enterprise templates and standards. This responsibility is centered in the Office of the DON CIO with the Enterprise Architecture and Standards Competency Unit.

6.1.2 Approval and Adoption by DON CIO Board of Representatives

This publication represents the consensus of the ITI IPT plus invited reviewers. It is now subject to review by additional representatives from across the DON prior to its submission to the DON CIO Board of Representatives for formal approval. It is critical that the Board approves the submitted framework and that each of the command CIOs agrees to apply the framework in their organizations. We must move as quickly as possible to using one common framework. If there are problems with the draft framework, we will fix them.

6.1.3 Acceptance and Adoption by Other Key Enterprise Projects

Although adoption by the IT community is essential, it is not sufficient. To be effective, the EAF must be accepted and adopted by the broader community of planners involved in RMA/RBA-related projects. This is particularly true of all projects that affect systems architectures or designs and all initiatives that are looking to improve acquisition and investment planning processes. IWARs is a case in point. A representative of IWARs participated in the framework workshop, but more must be done to ensure that the EAF is useful to this key transformation initiative.

6.1.4 Population & Maintenance of Enterprise Reference Models

As was stated earlier, the EAF provides the super-structure in which to position all architecture modeling initiatives. The framework really takes on meaning when it is populated with reference models for each of the architecture elements and their relationships. Recall that the UNTL is a reference model for the Mission Essential Task element of the Operational View. The good news is that the DON has already

developed and, in most cases, automated reference models for most of the elements in the EAF. These just need to be consolidated within a common toolset and repository. Space and Naval Systems Warfare Command (SPAWAR) has already created a compact disc version of the Naval Architecture Database V2.1G that has collected many of the existing architecture and standards in one place. It is important to align these many initiatives with the EAF and establish the means to make these truly enterprise-wide and to ensure that the reference models are kept up-to-date.

6.1.5 Provision of an Architecture Modeling Tool and Repository

The sheer volume and complexity of architecture models required to populate the framework and track the many inter-related plans dictate the need for a highly functional modeling tool and repository. Most of the existing models are being maintained in spreadsheets, file managers, drawing tools, and document processors. Architectures are, after all, primarily comprised of graphical and pictorial drawings with associated attributes and information. What is needed is a tool that allows the creation, maintenance, and browsing of architecture drawings across the entire EAF. This tool must be tightly linked to a model and information repository that permit the development and generation of item specifications, matrices, and reports plus provide effective search capabilities. The tool and repository must then be made available with appropriate controls to the various planning communities working on selected models.

The following provides a checklist for the desired features of the DON Architecture Modeling Tool and Repository:

General Capabilities:

- Model-based approach
- Object class orientation
- Ability to support all possible types of models in the EAF
- Ease of customization of models, classes and attributes to suit EAF, planning methodology, and user requirements
- Multi-media user interface
- Scaleable for enterprise-wide use
- Ability to link and import/export items from other tools

Modeling Capabilities:

- Definition and customization of the types of components and connectors, including visual appearance, labeling, object and diagram rules, and display options
- Assignment of diagram items to the object classes and repository entries (new or existing)
- Intelligent diagrams - maintenance and knowledge of connections and interconnections
- Multiple sub-views - ability to create drill-down capability and/or branching to any related model in context of the diagram and its items
- Population of attributes through formatted specification sheets, matrices, or imported fields

Viewing Capabilities:

- Model navigation approach leveraging multiple sub-views for diagrams and objects (in context)
- Attribute display by specification sheet and/or matrix

- Matrix population based on diagrams or repository
- Relationship display through connectors, drill-downs, and relationship matrices
- Linkage from diagrams to other related data bases, documents, and web sites (knowledge management environment)
- Model and class browsing capability
- Generation of “reports” and documents from the repository

Operational Capabilities:

- Central maintenance of EAF and reference models
- Project administration controls (models, classes, users)
- Multi-user project repository with write controls
- Ability to work offline and independent from master repository and project repository
- Project back-up and recovery support
- Tracing of user and change dates

6.1.6 Development of Model-based Planning and Design Methods

The availability of a well-populated framework with reference models, templates, and standards supported by a powerful modeling tool and repository sets the stage for revolutionizing the planning methods of the DON. Before being able to plan for RMA/RBA, we must first revolutionize our way of planning! By developing and using model-based planning methods rooted in the enterprise reference models, we make our planning processes both more effective and efficient. For example, there is current activity within the fleets using the elements in the Operational View to determine requirements for IT Infrastructure. Once the association of workstations (equipment) to personnel is established within operational areas, the requirement for the number of workstations and the related network connectivity can be automatically derived from headcount and location data in conjunction with templates. As these model-based planning methods are developed, they can be made available for use across the enterprise.

6.1.7 Establishment of Re-Use Center and Library for Components

Another way to leverage the investment in architecture management is to establish a re-use center to support the enterprise. This center should be responsible for logging, certifying, and distributing components for re-use. This can apply to component designs, but is extremely applicable to the software community. Component-based development methodologies are now becoming popular, especially as object-oriented designs and code become available. Major breakthroughs in development productivity are achievable through re-use, but it takes a well-organized approach to leverage this opportunity. The framework and reference models are a prerequisite to identifying and classifying reusable components.

6.2 Using the Framework for the ITI IPT

The EAF is being introduced in conjunction with the work of the ITI IPT. This IPT is the first of a number of enterprise IPTs to be stood-up by the DON CIO Board of Representatives to bring an enterprise-wide perspective to planning for the usage of information technologies and information management systems in supporting the transformation of the DON. The EAF, as described in the previous chapters, can now be used to position the work of the ITI IPT.

In addition to producing this EAF, the IPT was primarily tasked with developing architectures and planning templates to support enterprise-wide networking as well as addressing the related governance and requirements planning issues. Each of these taskings is presented below in the context of the EAF.

6.2.1 Enterprise Network Architecture Framework

The mandate of the ITI IPT was focused on developing architectures and templates related to enterprise networks. This scope is shown in Figure 6-1 as outlined in yellow.

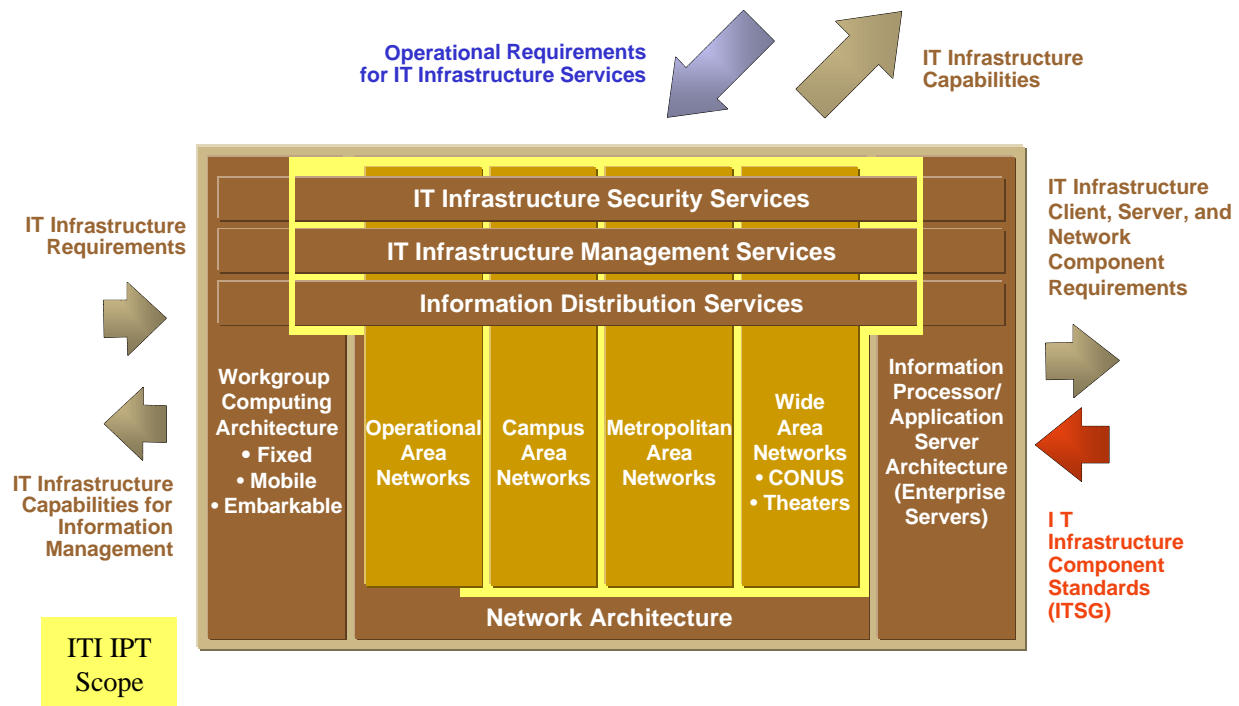


Figure 6-1. Information Technology Architecture Framework Showing Scope of ITI IPT

The Network Architecture View provides the framework for defining the components of the Enterprise Architecture. Note that the mandate included the Wide Area Network Architecture Plan and templates to address the common architecture designs for Metropolitan and Campus (or Base) Area Networks. The Operational Area Networks are also candidates for templates reflecting common architectures for platform-based networks, but were not included in the scope of the IPT. The same is true for Workgroup Computing Architectures and the Information Processor/Application Server Architectures.

A key focus of the IPT was on the common network services that must operate across the technology infrastructure. These address the three primary service areas of security, management, and information distribution. All of the IT architectures and the three service areas can be broken into their respective architecture elements. This model is shown in Figure 6-2.

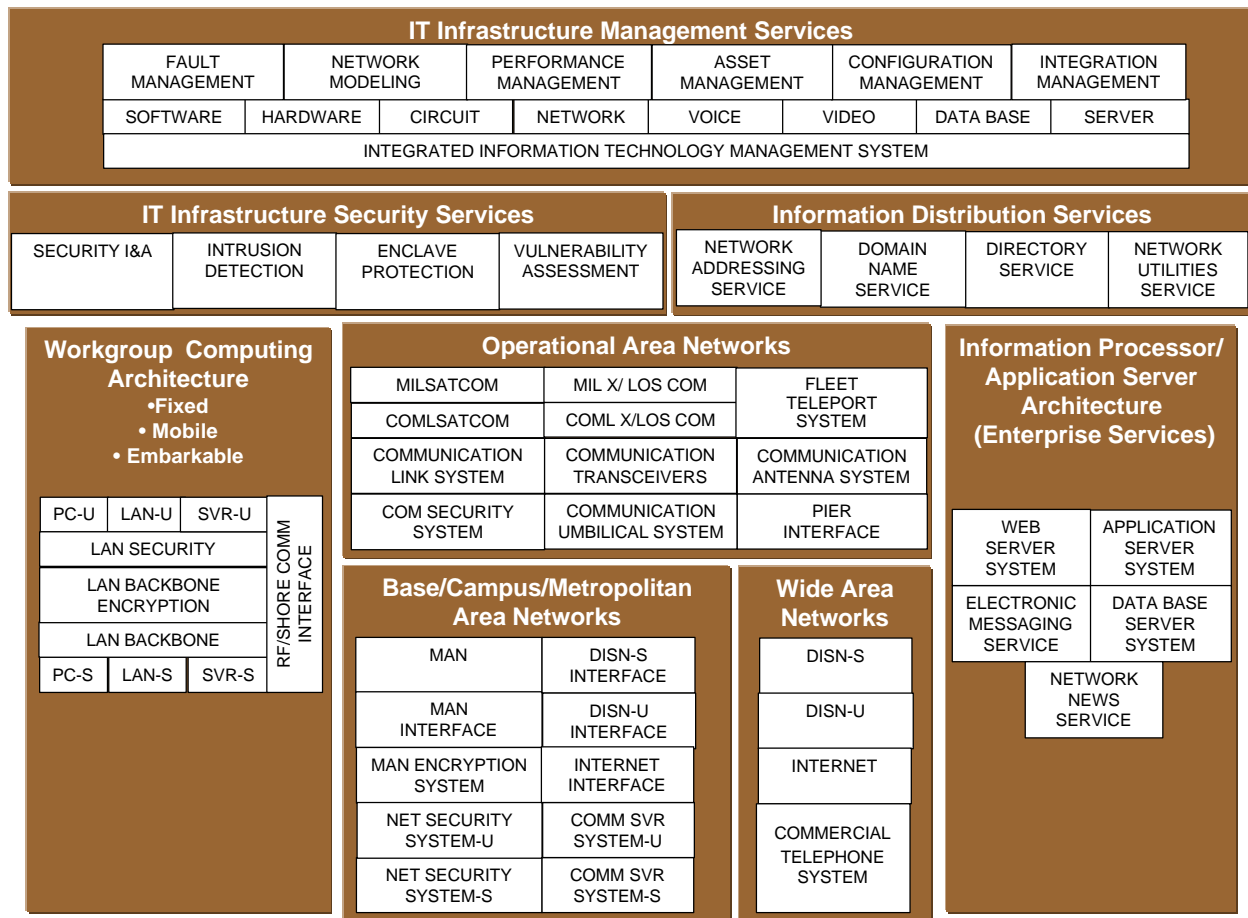


Figure 6-2. The Architecture Components of the DON ITI Architecture

This framework provides the context for identifying the specific connectivity requirements between components of the network architecture and the specific services that must operate across the network.

The major transformations to be achieved in network architecture are summarized in Figure 6-3.

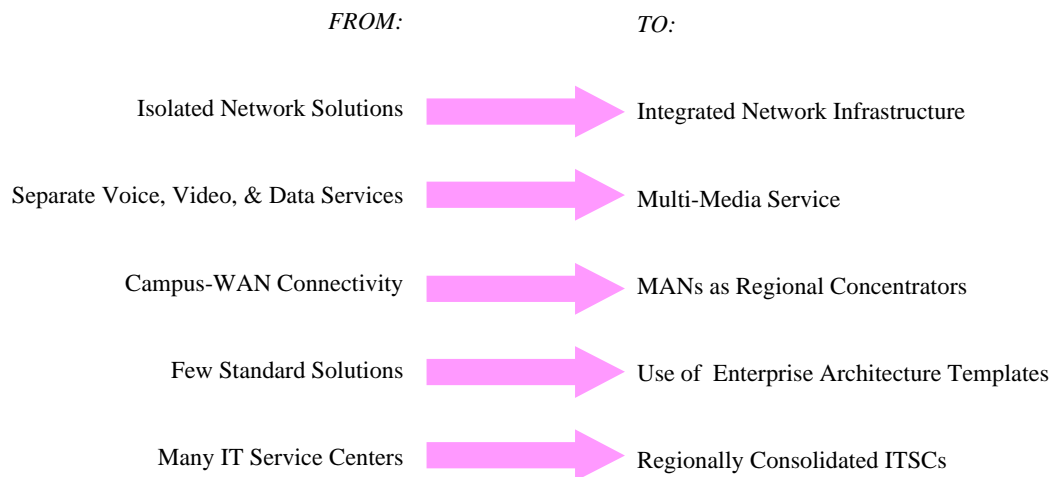


Figure 6-3. Desired Transformations in DON Network Architecture

6.2.2 ITI Governance Framework

The introduction of an Enterprise IT Infrastructure presupposes some major changes in the way IT planning, design, procurement, and operations are conducted throughout the DON. To develop meaningful recommendations for governance, the IPT developed a “business model” to reflect the functional activities and organizational considerations required to fulfill the mission of providing enterprise IT Services to DON Commands. This “business model” is an extension of the MET model in the Operational View of the framework. The UNTL is a well-developed hierarchical model of the military tasks that can now be supplemented with models of the business-related tasks to complete the enterprise MET model.

Figure 6-4 shows a model containing all of the functional components (i.e. METs) required to fulfill the mission of providing IT services for the DON. Note that this mission is an underlying support function for most, if not all, of the METs in the UNTL.

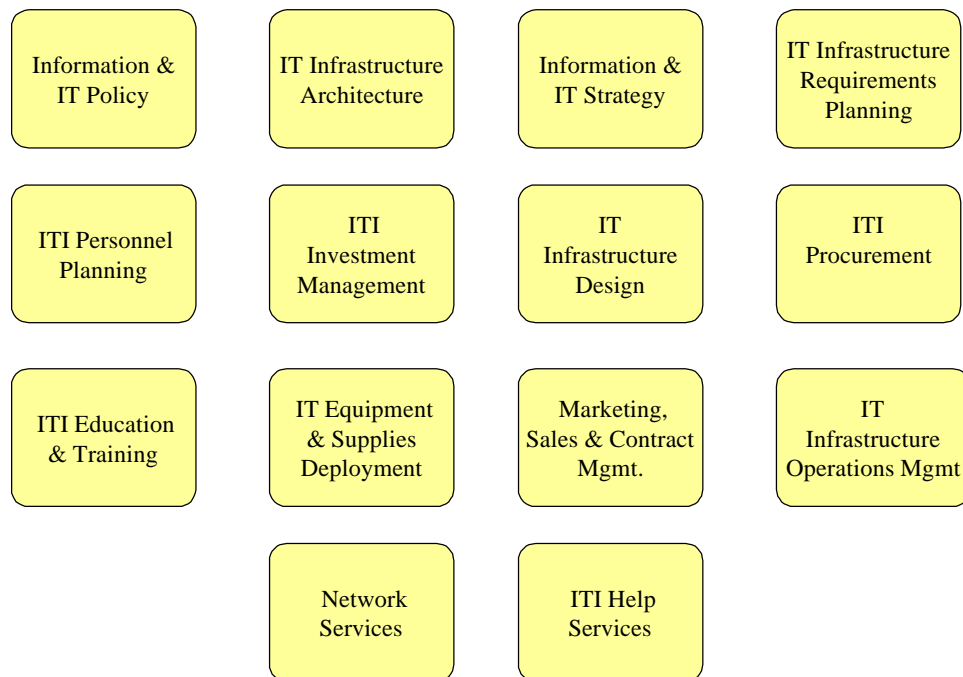


Figure 6-4. The Mission Essential Tasks for IT Infrastructure Services

The IPT makes use of this business model as a framework for assessing governance issues and presenting recommendations regarding realignment of organizational accountabilities and other operational changes required to effectively plan, design, procure, manage, and operate IT Infrastructure Services. See Volumes III and IV for presentation of these findings.

The required transformation in ITI Governance is summarized in Figure 6-5.

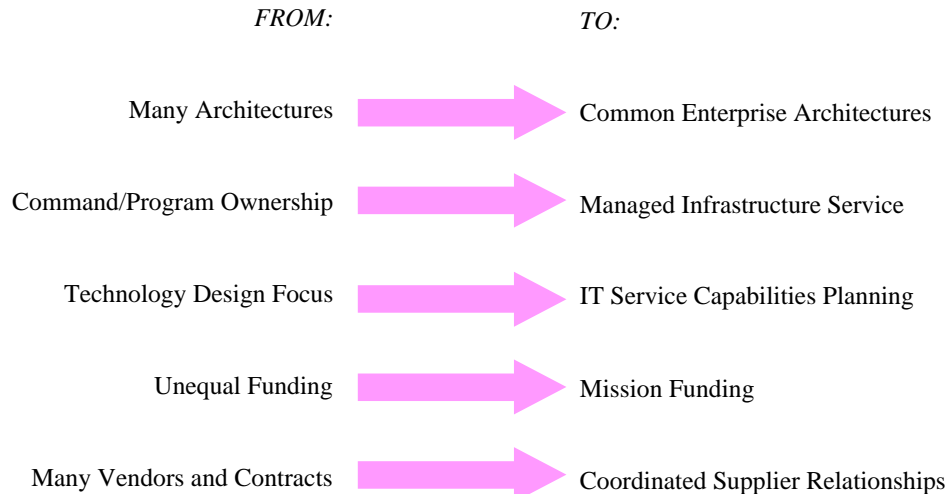


Figure 6-5. Desired Transformations in IT Infrastructure Governance

6.2.3 ITI Services Requirements Planning Framework

One of the METs in the IT Services Business Model (Figure 6-4) is IT Infrastructure Requirements Planning. The EAF is particularly well-suited to assist in transforming the way in which the DON plans for IT services because it was constructed with the primary goal of linking technology planning decisions to the strategic mission requirements of the DON.

There are two primary sources of IT Infrastructure Service requirements as shown in Figure 6-1. The first is the Operational View that provides for each organizational unit, whether fixed, mobile, or embarkable, the number and type of personnel requiring access to ITI Service Capabilities to perform their tasks. This easily leads to a determination of the number of “seats” by organization, by facility, and by platform that must be supported by the IT Infrastructure. By using the framework and linking IT planning to personnel planning and rules for allocation of workstations and services in the different operating environments, it is possible to directly forecast demand for access to the various ITI capabilities. This approach is now being used by the fleets and across N6.

The second source of requirements comes from the Information Systems Architecture. These relate to the demands for information and application processing capabilities. This primarily affects the size, numbers, and placement of application servers to meet the enterprise-wide requirements which, in turn, affects the traffic patterns on the network.

Volume IV contains process models that identify the recommended approach for performing ITI requirements planning. Three different types of processes are identified. The first relates to the functional requirements for the ITI Architecture; the second, to the functional requirements related to designing and developing the specific capabilities; and the third, to the functional requirements related to the operation of the ITI Services.

Figure 6-6 shows the major transformations that are expected in the processes for planning for IT Infrastructure Services.

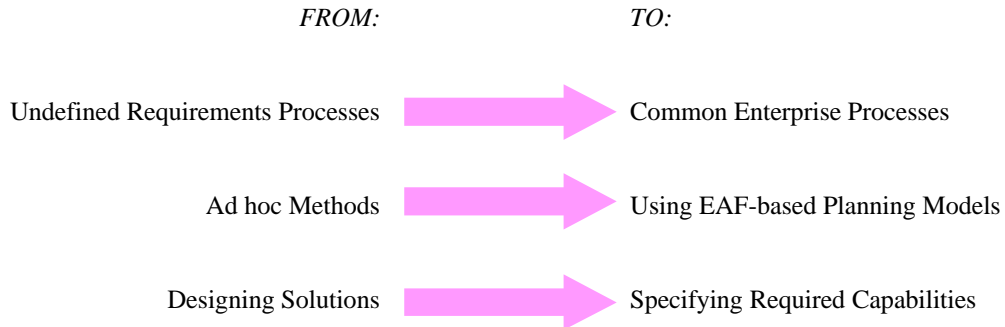


Figure 6-6. Desired Transformations in IT Infrastructure Requirements Planning

6.3 Conclusion

The transformation of the DON will provide many opportunities for leadership. A good leader will prepare for the challenge by adopting the best means available to ensure success. The opportunities for improving the enterprise and discovering innovative means of leveraging information and information technologies are tremendous; but so are the challenges.

Planning methods rooted in sound architecture principles and practices have proven to be the best means of tackling complex human and technical problems. The introduction of the EAF at this critical time of (r)evolution is meant to provide these leaders with an important management tool. The broader the adoption and application of these model-based planning methods, the greater the return to the DON.